arXiv:1704.06988v1 [stat.ME] 23 Apr 2017

Extended ensemble Kalman filters for high-dimensional hierarchical state-space models

Matthias Katzfuss* Jonathan R. Stroud† Christopher K. Wikle‡

Abstract

The ensemble Kalman filter (EnKF) is a computational technique for approximate inference on the state vector in spatio-temporal state-space models. It has been successfully used in many real-world nonlinear data-assimilation problems with very high dimensions, such as weather forecasting. However, the EnKF is most appropriate for additive Gaussian state-space models with linear observation equation and without unknown parameters. Here, we consider a broader class of hierarchical state-space models, which include two additional layers: The parameter layer allows handling of unknown variables that cannot be easily included in the state vector, while the transformation layer can be used to model non-Gaussian observations. For Bayesian inference in such hierarchical state-space models, we propose a general class of extended EnKFs, which approximate inference on the state vector in suitable existing Bayesian inference techniques (e.g., Gibbs sampler or particle filter) using the EnKF or the related ensemble Kalman smoother. Extended EnKFs enable approximate, computationally feasible filtering and smoothing in many high-dimensional, nonlinear, and non-Gaussian spatio-temporal models with unknown parameters. We highlight several interesting examples, including assimilation of heavy-tailed and discrete data, and filtering and smoothing inference on model parameters.

Keywords: data assimilation; spatio-temporal statistics; geoscience applications; ensemble Kalman smoother; particle filter; online inference

*Texas A&M University (katzfuss@tamu.edu)
†Georgetown University
‡University of Missouri

1 Introduction

1.1 Inference in state-space models

A state-space model (SSM) is a very general representation of a system observed sequentially over time. A SSM describes the temporal evolution of the system state, and the relationship of the state to observations. Here, we consider spatio-temporal SSMs describing the evolution of one or more (discretized) spatial fields over time, which often arise in geoscience applications. Such SSMs typically exhibit two features that make inference challenging: nonlinear evolution and high state dimension. While the Kalman filter (Kalman, 1960) provides exact inference for linear Gaussian SSMs, it is not suited for nonlinear or high-dimensional SSMs.

For nonlinear or non-Gaussian SSMs, inference is often carried out using sequential Monte Carlo (SMC) methods (e.g., Doucet et al., 2001). SMC methods approximate distributions at each time point by a ensemble representation, which can be easily propagated forward over time according to the true state evolution. However, most SMC methods update the ensemble based on new data using a reweighting or resampling step (e.g., Gordon et al., 1993). It is well known that such methods, which we will collectively refer to as particle filters here, suffer from weight degeneracy (i.e., all but one weight become essentially zero) in high-dimensional problems (e.g., Snyder et al., 2008), although some recent results suggest that it might in principle be possible to overcome this curse of dimensionality (e.g., Rebeschini and Van Handel, 2015).

An alternative SMC method called the ensemble Kalman filter (EnKF) was proposed by Evensen (1994) in the geophysics literature (see Houtekamer and Zhang, 2016, and Katzfuss et al., 2016, for reviews from an atmospheric-science and a statistical perspective, respectively). The EnKF also uses an ensemble representation that can be propagated forward in the same way as in the particle filter, but instead of reweighting, the EnKF updates the ensemble based on data using a linear "shift" that approximates the best linear update. While this update generally introduces some bias, it has also allowed the EnKF to be successfully applied to nonlinear spatio-temporal SSMs with millions of dimensions (e.g., Szunyogh et al., 2008; Whitaker et al., 2008; Houtekamer et al., 2014).

The EnKF is most appropriate for fully specified state-space models in which the data are approximately Gaussian and are linearly related to the state. The linear shifting update in the EnKF can also be highly suboptimal for non-Gaussian observations. The most popular method for dealing with unknown parameters in the EnKF framework is state augmentation (Anderson, 2001), which refers to simply including the parameters in the state vector. However, the resulting linear update can be highly problematic for several types of parameters (e.g., variance parameters; DelSole and Yang, 2010). Thus, for the purpose of this article, we use the term parameters for any (static or time-varying) unknown variables that cannot be simply included in the state, because the linear EnKF update would not be appropriate. In contrast, the state vector will typically consists of one or more (discretized) spatial fields.

1.2 Hierarchical state-space models

We consider here a more general class of *hierarchical state-space models* (HSSMs), which adds two layers to the observation and evolution layers in a "standard" SSM (cf. Cressie and Wikle, 2011, Ch. 7). The first additional layer, termed the *transformation layer*, allows the actual measurements to be non-Gaussian, as long as they can be expressed as some transformation of

latent Gaussian variables. Second, the *parameter layer* describes the evolution model for unknown parameters, which are treated as random. Here, we consider as parameters all unknown quantities that cannot be handled well using the state-augmentation approach described above. This class of HSSMs is very general and can be used to describe most spatio-temporal systems observed sequentially over time, including unknown parameters in any part of the model, non-Gaussian and non-linear observations, mixture models, and even higher-order Markov models (by extending the state-space). Some interesting examples will be given in Section 5.

More specifically, for discrete time points $t = 1, 2, \ldots$, our HSSM is given by

$$\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t \sim f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t), \tag{1}$$

$$\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}_t \sim \mathcal{N}_{m_t}(\mathbf{H}_t\mathbf{x}_t, \mathbf{R}_t), \tag{2}$$

$$\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\theta}_t \sim \mathcal{N}_n(\widetilde{\mathbf{x}}_t = \mathcal{M}_t(\mathbf{x}_{t-1}), \mathbf{Q}_t), \tag{3}$$

$$\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1} \sim p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}), \tag{4}$$

where $\mathbf{z}_t$ is the m_t-dimensional vector of measurements at time t, $\mathbf{y}_t$ is a corresponding m_t-dimensional "latent observation" vector, $\mathbf{x}_t$ is the n-dimensional unobserved state vector of interest, and the initial state is $\mathbf{x}_0|\boldsymbol{\theta}_0 \sim \mathcal{N}_n(\boldsymbol{\mu}_{0|0}, \boldsymbol{\Sigma}_{0|0})$ with $\boldsymbol{\theta}_0 \sim p_0(\boldsymbol{\theta}_0)$. The quantities $\mathbf{H}_t$, $\mathcal{M}_t$, $\mathbf{R}_t$, $\mathbf{Q}_t$, and f_t may all depend on the parameter vector $\boldsymbol{\theta}_t$. Sometimes we make the dependence explicit by writing $\mathbf{H}_t(\boldsymbol{\theta}_t)$ (and so forth). The $\mathbf{z}_t$ and $\mathbf{y}_t$ are serially independent given the states and parameters.

A standard (Gaussian) SSM is given by only the observation and evolution layer in (2) and (3), respectively. We have added to this the transformation layer (1) and the parameter layer (4). Note that our HSSM reduces to a standard Gaussian SSM when the latent observations $\mathbf{y}_t$ and the parameters $\boldsymbol{\theta}_t$ are known, or equivalently, if $f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t) = \delta_{\mathbf{y}_t}(\mathbf{z}_t)$ is a is a point mass at $\mathbf{y}_t$, and $p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1})$ is a point mass at the known value of $\boldsymbol{\theta}_t$. In this sense, our HSSM could also be described as a *conditionally latent-Gaussian SSM*.

We will be concerned with situations in which n and or m_t are very large (and usually $n > m_t$), and $\mathbf{x}_t$ is typically a discretized version of one or more spatial fields. Further, we view $\mathcal{M}_t$ as a "black box" evolution operator that might be nonlinear and intractable, in the sense of being expensive to evaluate or not being available in closed form.

1.3 Extended ensemble Kalman filters

For HSSMs of the form (1)–(4), the goal here is to make (approximate) Bayesian inference on the state $\mathbf{x}_t$ and the parameters $\boldsymbol{\theta}_t$. Throughout, let $[x|z]$ denote the conditional distribution (i.e., probability density or mass function) of some random variable x given z. We consider both the filtering perspective, which at each time t aims to find the filtering distribution $[\mathbf{x}_t, \boldsymbol{\theta}_t|\mathbf{z}_{1:t}]$ given the data $\mathbf{z}_{1:t} := \{\mathbf{z}_1, \ldots, \mathbf{z}_t\}$ collected up to time t, and the smoothing perspective, which aims to obtain the posterior distribution $[\mathbf{x}_t, \boldsymbol{\theta}_t|\mathbf{z}_{1:T}]$ for all t within a specific time period $\{1, \ldots, T\}$. Our focus is mostly on time-varying parameters, but static parameters (i.e., $\boldsymbol{\theta}_0 = \boldsymbol{\theta}_1 = \ldots$) are also discussed.

We propose a *simple general idea* for such inference in a HSSM: Take a suitable existing Bayesian inference technique (e.g., a particle filter or a Gibbs sampler), but approximate the part of the algorithm that requires sampling from or integrating out the high-dimensional state vector $\mathbf{x}_t$ with an EnKF or with the related ensemble Kalman smoother (EnKS). This leads to a general class of *extended ensemble Kalman filters and smoothers*, which can be applied

		Filtering/smoothing methods		
Target distrib.	Dim. of $\boldsymbol{\theta}_t$	Linear Gaussian	Particle	Ensemble
$[\mathbf{x}_t \vert \boldsymbol{\theta}_{1:t}, \mathbf{y}_{1:t}]$, $\forall t$	($\boldsymbol{\theta}$ known)	Kalman filter	Particle filter	EnKF
$[\mathbf{x}_{1:T} \vert \boldsymbol{\theta}_{1:T}, \mathbf{y}_{1:T}]$	($\boldsymbol{\theta}$ known)	Kalman smoother	Particle smoother	EnKS (Alg. 1)
$[\mathbf{x}_t, \boldsymbol{\theta}_t \vert \mathbf{z}_{1:t}]$, $\forall t$	high low	RWFFBS RBPF	Online MCMC EARBPF	**GEnKF (Alg. 2)** **PEnKF (Alg. 3)**
$[\mathbf{x}_{1:T}, \boldsymbol{\theta}_{1:T} \vert \mathbf{z}_{1:T}]$	high low	Gibbs with FFBS MH with KF	Particle Gibbs PMMH	**GEnKS (Alg. 4)** **MHEnKF/PEnKS (Alg. 5)**

Table 1: Existing and our new (in bold) methods for filtering and smoothing tasks for known and unknown parameters. The algorithms for high-dimensional $\boldsymbol{\theta}_t$ generally require its full-conditional distribution to be available in closed form. "Linear Gaussian" contains methods only suitable for linear Gaussian SSMs (conditional on $\boldsymbol{\theta}_t$). The methods in the column "Particle" do not scale to large state dimension n. Abbreviation and main references: KF = Kalman filter (Kalman, 1960); particle filter (Gordon et al., 1993); EnKF = ensemble Kalman filter (Evensen, 1994); Kalman smoother (Rauch et al., 1965); particle smoother (Kitagawa, 1996); EnKS = ensemble Kalman smoother (Evensen and van Leeuwen, 2000); RWFFBS = rolling-window FFBS (Polson et al., 2008); Online MCMC (Berzuini et al., 1997); RBPF = Rao-Blackwellized PF (Doucet et al., 2000a); EARBPF = exact approximate RBPF (Johansen et al., 2012); FFBS = forward-filtering backward-sampling (Carter and Kohn, 1994); particle Gibbs (Andrieu et al., 2010); MH with KF = Metropolis-Hastings based on KF likelihood (Schweppe, 1965); PMMH = particle marginal Metropolis-Hastings (Andrieu et al., 2010)

to HSSMs. The key is that when conditioning on the $\mathbf{y}_t$ and $\boldsymbol{\theta}_t$, our HSSM reduces to a standard Gaussian SSM with potentially nonlinear evolution, for which the EnKF is often well suited. Table 1 puts our proposed algorithms in the context of existing particle approaches or approaches for linear Gaussian SSMs.

Our new extended EnKF methods can be viewed from two perspectives: For readers familiar with the EnKF, our extended EnKFs greatly expand the class of problems to which the EnKF can be applied. For readers more familiar with established Bayesian inference methods, such as particle filtering or Gibbs sampling, our extended EnKF methods allow similar (but approximate) inference in higher-dimensional SSMs than the standard methods.

In the EnKF literature, several approaches exist for estimation of specific tuning parameters, such as inflation and localization parameters (e.g., Wang and Bishop, 2003; Anderson, 2007a), including a so-called hierarchical ensemble filter (Anderson, 2007b) that uses an ensemble of ensembles (similarly to our particle EnKF described later) to estimate a localization function. Here we focus instead on inference on model parameters that explicitly appear in the HSSM in (1)–(4), and we regard the EnKF tuning parameters as known.

As the EnKF is an approximate technique for inference on the state vector, our extended EnKF techniques also can only produce approximate inference. However, in many complex problems (e.g., when the state dimension is high), exact inference is impossible, and approximate techniques such as approximate Bayesian computation (ABC; Sisson et al., 2007; Marin et al., 2012), synthetic likelihood (e.g., Wood, 2010) and integrated nested Laplace approximations (INLA; Rue et al., 2009; Lindgren et al., 2011) have generally gained in popularity. While so-called exact approximate approaches (e.g., Andrieu et al., 2010; Johansen et al., 2012) result in exact inference given infinite computational resources, we will show here numerically that, in finite computation time, our extended EnKF techniques can outperform exact approximate techniques even for moderately high dimensions.

1.4 Overview and contributions

This article is organized as follows. In Section 2, we review some existing work on the EnKF and related issues from a statistical perspective. In Section 3, we provide preliminaries to lay the foundation for the development of our extended EnKF methods, including new insights on likelihood approximations in high-dimensional SSMs. In Section 4, we derive and present our extended ensemble Kalman filters and smoothers. In Section 5, we consider several examples of HSSMs, and we offer numerical comparisons of our methods to existing ones. We conclude in Section 6.

2 Review of the ensemble Kalman filter (EnKF)

We begin with a review of the ensemble Kalman filter (EnKF) in a form that will be useful for the extensions presented later. The EnKF is most appropriate for Gaussian SSMs as in (2)–(3) for which $\mathbf{y}_t$ and $\boldsymbol{\theta}_t$ are known. This is essentially a special case of our general HSSM for known parameters and linear Gaussian observations, where $f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t)$ is a point mass at $\mathbf{y}_t$. Hence, assume temporarily the SSM (2)–(3) with all $\mathbf{y}_t$ and $\boldsymbol{\theta}_t$ known (and we will often suppress explicit conditioning on $\boldsymbol{\theta}_{1:t}$ for notational simplicity).

As described in Section 1.2, we assume here that the evolution operator $\mathcal{M}_t$ is intractable or computationally expensive, and so all we are able to do with $\mathcal{M}_t$ is to "apply" it to a finite (and typically rather small) number of state vectors. Thus, sequential inference over time must be based on an *ensemble* of state vectors that represents (i.e., approximates) the relevant distributions at each time point.

More specifically, assume that the *filtering ensemble* at time $t-1$ is given by $\mathbf{x}_{t-1|t-1}^{(1:N)} := \{\mathbf{x}_{t-1|t-1}^{(1)}, \ldots, \mathbf{x}_{t-1|t-1}^{(N)}\}$, meaning that the filtering distribution at time $t-1$ is approximated as

$$[\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}] = \tfrac{1}{N}\textstyle\sum_{j=1}^{N} \delta_{(\mathbf{x}_{t-1|t-1}^{(j)})}(\mathbf{x}_{t-1}), \tag{5}$$

where δ is the Dirac delta (point mass) function. The *forecast ensemble* is obtained by applying the true evolution model to each member of the filtering ensemble:

$$\widetilde{\mathbf{x}}_{t|t-1}^{(j)} = \mathcal{M}_t(\mathbf{x}_{t-1|t-1}^{(j)}), \qquad j = 1, \ldots, N, \tag{6}$$

resulting in the forecast ensemble $\widetilde{\mathbf{x}}_{t|t-1}^{(1:N)} = \{\widetilde{\mathbf{x}}_{t|t-1}^{(1)}, \ldots, \widetilde{\mathbf{x}}_{t|t-1}^{(N)}\}$.

2.1 Updating the state

Based on new data $\mathbf{y}_t$ observed at time t, we would then like to obtain or approximate the filtering distribution of the state vector, which is the posterior distribution of $\mathbf{x}_t$ given $\mathbf{y}_{1:t}$. Hence the task is to update the forecast ensemble in (6) to obtain a filtering ensemble that takes into account the information in $\mathbf{y}_t$.

Based on (5), the prior distribution at time t would in principle be given by

$$[\mathbf{x}_t|\mathbf{y}_{1:t-1}] = \textstyle\int [\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{y}_{1:t-1}] dP(\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}) = \tfrac{1}{N}\sum_{j=1}^{N} \mathcal{N}_n(\mathbf{x}_t|\widetilde{\mathbf{x}}_{t|t-1}^{(j)}, \mathbf{Q}_t), \tag{7}$$

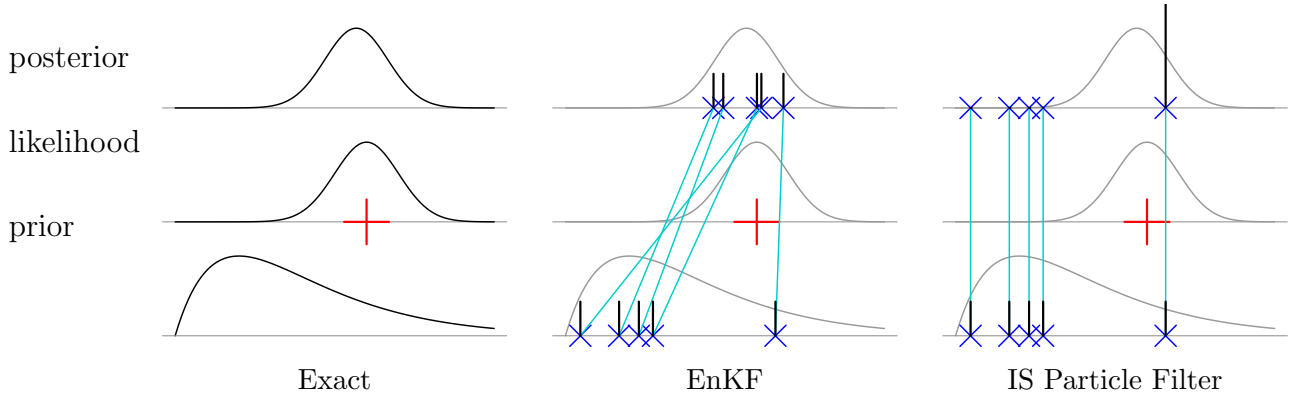

Figure 1: Comparison of updates in a simple one-dimensional example with a non-Gaussian prior distribution. The bars above the ensemble/particles are proportional to the weights. Both approximation methods start with the same, equally weighted $N = 5$ prior samples. Even for this one-dimensional example, the importance sampling (IS) weights degenerate, while the EnKF shifting update obtains a better representation of the posterior distribution.

where the integral is with respect to the measure induced by the distribution of $\mathbf{x}_{t-1}|\mathbf{y}_{1:t-1}$. However, based on the mixture prior distribution in (7), the filtering distribution updated using the new observation $\mathbf{y}_t$ would have the form

$$\begin{aligned}[\mathbf{x}_t|\mathbf{y}_{1:t}] &\propto [\mathbf{y}_t|\mathbf{x}_t, \mathbf{y}_{1:t-1}][\mathbf{x}_t|\mathbf{y}_{1:t-1}] = \tfrac{1}{N}\textstyle\sum_{j=1}^{N} \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\mathbf{x}_t, \mathbf{R}_t)\mathcal{N}_n(\mathbf{x}_t|\widetilde{\mathbf{x}}_{t|t-1}^{(j)}, \mathbf{Q}_t) \\ &\propto \textstyle\sum_{j=1}^{N} \alpha_t^{(j)}[\mathbf{x}_t|\mathbf{y}_t, \mathbf{x}_{t-1|t-1}^{(j)}], \end{aligned} \tag{8}$$

with the filtering or posterior weights

$$\alpha_t^{(j)} = [\mathbf{x}_{t-1|t-1}^{(j)}|\mathbf{y}_{1:t}] \propto \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\widetilde{\mathbf{x}}_{t|t-1}^{(j)}, \mathbf{H}_t\mathbf{Q}_t\mathbf{H}_t' + \mathbf{R}_t). \tag{9}$$

This reweighting of the jth forecast ensemble member $\widetilde{\mathbf{x}}_{t|t-1}^{(j)}$ from prior weight $1/N$ to posterior weight $\alpha_t^{(j)}$ is essentially what happens in an importance sampler or in the update step of a particle filter that integrates out the innovation error with covariance $\mathbf{Q}_t$. This reweighting update is illustrated in a toy example in Figure 1 for the simple case $m_t = 1$, $\mathbf{H}_t = 1$, and $\mathbf{Q}_t = \mathbf{0}$. However, as shown in Snyder et al. (2008) by setting their observation error covariance matrix $\mathbf{R}$ to be $\mathbf{H}_t\mathbf{Q}_t\mathbf{H}_t' + \mathbf{R}_t$ in our notation, the particle weights degenerate, meaning that one mixture component (i.e., one ensemble member) would get (nearly) all of the weight, unless the ensemble size N increases exponentially with the effective dimension given by the variance of the logarithm of the particle likelihood (see Section 3.2.1 below). Thus, in our high-dimensional setting where only a small ensemble size N is computationally affordable, this update based on reweighting is problematic.

2.1.1 The filtering distribution for linear evolution

In contrast, the EnKF update does not reweight but rather shifts the forecast ensemble. The form of the shift is motivated by the fact that if the evolution operators $\mathcal{M}_t$ are all linear, then the true prior distribution will be a multivariate Gaussian distribution. In that case, it is natural to approximate the mixture of Gaussians in the prior distribution (7) by a (non-mixture) Gaussian distribution:

$$[\mathbf{x}_t|\mathbf{y}_{1:t-1}] = \mathcal{N}_n(\mathbf{x}_t|\widehat{\boldsymbol{\mu}}_{t|t-1}, \widehat{\boldsymbol{\Sigma}}_{t|t-1}), \tag{10}$$

where $\widehat{\boldsymbol{\mu}}_{t|t-1}$ and $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$ are estimates of the prior mean and covariance matrix (see Section 2.2 below) based on the forecast ensemble obtained in (6). Given this multivariate Gaussian prior distribution, the posterior (filtering) distribution would also be multivariate Gaussian:

$$[\mathbf{x}_t|\mathbf{y}_{1:t}] \propto [\mathbf{y}_t|\mathbf{x}_t, \mathbf{y}_{1:t-1}][\mathbf{x}_t|\mathbf{y}_{1:t-1}] = \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\mathbf{x}_t, \mathbf{R}_t)\mathcal{N}_n(\mathbf{x}_t|\widehat{\boldsymbol{\mu}}_{t|t-1}, \widehat{\boldsymbol{\Sigma}}_{t|t-1}) \propto \mathcal{N}_n(\mathbf{x}_t|\widehat{\boldsymbol{\mu}}_{t|t}, \widehat{\boldsymbol{\Sigma}}_{t|t}),$$

where $\widehat{\boldsymbol{\mu}}_{t|t} = \widehat{\boldsymbol{\mu}}_{t|t-1} + \widehat{\mathbf{K}}_t(\mathbf{y}_t - \mathbf{H}_t\widehat{\boldsymbol{\mu}}_{t|t-1})$, $\widehat{\boldsymbol{\Sigma}}_{t|t} = (\mathbf{I} - \widehat{\mathbf{K}}_t\mathbf{H}_t)\widehat{\boldsymbol{\Sigma}}_{t|t-1}$, and

$$\widehat{\mathbf{K}}_t = \widehat{\boldsymbol{\Sigma}}_{t|t-1}\mathbf{H}_t'(\mathbf{H}_t\widehat{\boldsymbol{\Sigma}}_{t|t-1}\mathbf{H}_t' + \mathbf{R}_t)^{-1} \tag{11}$$

is an estimate of the Kalman gain, $\mathbf{K}_t$. A sample from this posterior distribution can be obtained via shifting the forecast ensemble members by essentially carrying out conditional simulation (e.g., Katzfuss et al., 2016):

$$\mathbf{x}_{t|t}^{(j)} = \mathbf{x}_{t|t-1}^{(j)} + \widehat{\mathbf{K}}_t(\mathbf{y}_t - \widetilde{\mathbf{y}}_t^{(j)}) =: \mathcal{U}_t(\widetilde{\mathbf{x}}_{t|t-1}^{(j)}|\mathbf{y}_t, \widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}, \boldsymbol{\theta}_t) \tag{12}$$

where $\mathbf{x}_{t|t-1}^{(j)} = \widetilde{\mathbf{x}}_{t|t-1}^{(j)} + \mathbf{w}_t^{(j)}$ is a prior ensemble member with $\mathbf{w}_t^{(j)} \overset{iid}{\sim} \mathcal{N}_n(\mathbf{0}, \mathbf{Q}_t)$, $\widetilde{\mathbf{y}}_t^{(j)} = \mathbf{H}_t\mathbf{x}_{t|t-1}^{(j)} + \mathbf{v}_t^{(j)}$ is a pseudo-observation with $\mathbf{v}_t^{(j)} \overset{iid}{\sim} \mathcal{N}(\mathbf{0}, \mathbf{R}_t)$, the forecast ensemble $\widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}$ enters the shifting update $\mathcal{U}_t$ through the estimated Kalman gain $\widehat{\mathbf{K}}_t$, and $\mathbf{H}_t$, $\mathbf{Q}_t$, and $\mathbf{R}_t$ are evaluated at the parameter value $\boldsymbol{\theta}_t$ (currently still assumed to be known).

2.1.2 The EnKF update

The (stochastic) EnKF update is exactly given by (12). This shifting update is derived based on assuming $\mathcal{M}_t$ is linear (and hence the prior distribution is multivariate Gaussian), and the resulting filtering distribution converges to the true filtering distribution as $N \to \infty$ (assuming the estimator $\boldsymbol{\Sigma}_{t|t-1}$ converges to the true forecast covariance; see Section 2.2 below). For finite N, in contrast to many other SMC methods, the EnKF update does not deteriorate for large n, as long as a reasonable estimate of $\mathbf{K}_t$ or $\boldsymbol{\Sigma}_{t|t-1}$ can be found. It has been shown repeatedly (e.g., Lei et al., 2010) that this shifting update also works surprisingly well for nonlinear evolution or non-Gaussian priors (see Figure 1 for an illustration), including operational, real-world applications (e.g., Houtekamer and Mitchell, 2005; Bonavita et al., 2010; Houtekamer et al., 2014). In fact, the shifting update can be viewed as an approximation (because $\mathbf{K}_t$ is estimated by $\widehat{\mathbf{K}}_t$) of the optimal linear update regardless of the true distribution (e.g. Le Gland et al., 2011; Law et al., 2016). This is similar to the notions in so-called linear Bayes estimation that considers linear estimators given only the first and second moments of the prior distribution and likelihood (e.g, Hartigan, 1969; Goldstein and Wooff, 2007). In fact, there is a close connection between ABC and linear Bayes estimation (e.g., Nott et al., 2014) and the EnKF can be shown to be equivalent to special cases of ABC algorithms (Nott et al., 2012).

2.2 Estimation of matrices

The EnKF requires estimates of the mean $\widehat{\boldsymbol{\mu}}_{t|t-1}$ and covariance matrix $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$ of the prior distribution in (10). The prior covariance is required for the EnKF update through the Kalman gain (11), and both mean and covariance are needed in the EnKF likelihood in Section 3.2.2 below.

The prior mean can simply be estimated by the sample mean of the forecast ensemble; that is, $\widehat{\boldsymbol{\mu}}_{t|t-1} = \frac{1}{N}\sum_{j=1}^{N} \widetilde{\mathbf{x}}_{t|t-1}^{(j)}$.

Estimation of the $n \times n$ prior covariance matrix based on small ensembles is often challenging in applications with large state dimension n, and regularization becomes necessary to avoid rank deficiency and spurious correlations. If some ordering or distance between elements of the state vector can be specified, such that the dependence between two elements is expected to decrease with increasing distance (e.g., when the state is a discretized version of a spatial field), a popular and effective regularization method is tapering (e.g., Furrer et al., 2006; Anderson, 2007b; Furrer and Bengtsson, 2007), which refers to Hadamard (entrywise) multiplication with a sparse positive definite correlation matrix defined based on the distances between the elements. This leads to the expression $\widehat{\boldsymbol{\Sigma}}_{t|t-1} = (\widetilde{\boldsymbol{\Sigma}}_{t|t-1} \circ \mathcal{T}_t) + \mathbf{Q}_t$, where $\widetilde{\boldsymbol{\Sigma}}_{t|t-1}$ is the sample covariance matrix of $\{\widetilde{\mathbf{x}}_{t|t-1}^{(1)}, \ldots, \widetilde{\mathbf{x}}_{t|t-1}^{(N)}\}$, and $\mathcal{T}_t$ is the sparse correlation matrix. Other regularization methods for states that are spatially referenced (e.g., Ott et al., 2004; Hunt et al., 2007) or not (e.g., Pourahmadi, 2013) will not be discussed here, but could also be applied in some of our methods described later.

2.3 Ensemble Kalman smoother

To approximate the smoothing distribution $[\mathbf{x}_t|\mathbf{y}_{1:T}]$ for some T, various extensions of the EnKF to an ensemble Kalman smoother have been proposed (van Leeuwen and Evensen, 1996; Evensen and van Leeuwen, 2000; Khare et al., 2008; Stroud et al., 2010; Bocquet and Sakov, 2014). The most widely-used smoother is the EnKS algorithm of Evensen and van Leeuwen (2000). Like the EnKF, the EnKS is a forward-only algorithm that is applicable even for intractable evolution operators $\mathcal{M}_t$. At each time t, it essentially carries out an EnKF update on the augmented state that includes the entire history, $\mathbf{x}_{1:t} = (\mathbf{x}_1', \ldots, \mathbf{x}_t')'$. This leads to the following algorithm:

Algorithm 1: Ensemble Kalman smoother (EnKS)
Start with an initial ensemble $\mathbf{x}_{0|0}^{(1)}, \ldots, \mathbf{x}_{0|0}^{(N)} \sim \mathcal{N}(\boldsymbol{\mu}_{0|0}, \boldsymbol{\Sigma}_{0|0})$. Then, at each time $t = 1, \ldots, T$, do the following for each $j = 1, \ldots, N$:

1. **Forecast Step:** Compute $\mathbf{x}_{t|t-1}^{(j)} = \mathcal{M}_t(\mathbf{x}_{t-1|t-1}^{(j)}) + \mathbf{w}_t^{(j)}$, where $\mathbf{w}_t^{(j)} \sim \mathcal{N}_n(\mathbf{0}, \mathbf{Q}_t)$.

2. **Update Step:** For $l = 1, \ldots, t$ compute $\mathbf{x}_{l|t}^{(j)} = \mathbf{x}_{l|t-1}^{(j)} + \widehat{\mathbf{K}}_{l,t}(\mathbf{y}_t - \widetilde{\mathbf{y}}_t^{(j)})$ where $\widetilde{\mathbf{y}}_t^{(j)} = \mathbf{H}_t\mathbf{x}_{t|t-1}^{(j)} + \mathbf{v}_t^{(j)}$ is the pseudo-observation with $\mathbf{v}_t^{(j)} \sim \mathcal{N}_{m_t}(\mathbf{0}, \mathbf{R}_t)$, $\widehat{\mathbf{K}}_{l,t} = \widehat{\boldsymbol{\Sigma}}_{l,t|t-1}\mathbf{H}_t'(\mathbf{H}_t\widehat{\boldsymbol{\Sigma}}_{t,t|t-1}\mathbf{H}_t' + \mathbf{R}_t)^{-1}$, and $\widehat{\boldsymbol{\Sigma}}_{l,t|t-1}$ is (a possibly regularized version of) the sample cross-covariance matrix of $\mathbf{x}_{l|t-1}^{(1:N)}$ and $\mathbf{x}_{t|t-1}^{(1:N)}$.

The resulting ensemble $\mathbf{x}_{t|T}^{(1:N)}$ is then an approximate sample from the smoothing distribution $[\mathbf{x}_t|\mathbf{y}_{1:T}]$, and more generally, $\mathbf{x}_{1:T|T}^{(1:N)}$ is an approximate sample from $[\mathbf{x}_{1:T}|\mathbf{y}_{1:T}]$

Note that regularization often results in $\widehat{\boldsymbol{\Sigma}}_{l,t|t-1} = \mathbf{0}$ for $l = 1, \ldots, t-k$ for some lag k. The computations in Step 2 of Algorithm 1 then only have to be carried out for k lags $l = t-k+1, \ldots, t$.

3 Preliminaries

In the previous section, we assumed that the latent $\mathbf{y}_{1:t}$ and the parameters $\boldsymbol{\theta}_{1:t}$ are all known. To obtain the necessary ingredients for the extended EnKF methods to be introduced in Section 4, we now extend the approaches reviewed in Section 2 to the more general HSSMs of the form (1)–(4) with non-Gaussian observations or unknown parameters.

Assume that we have a sample $\boldsymbol{\theta}_{t-1}^{(1)}, \ldots, \boldsymbol{\theta}_{t-1}^{(M)}$ from the filtering distribution of the parameters at time $t-1$. Except for special cases (see Section 3.1 below), we now need an ensemble of states, $\mathbf{x}_{t-1|t-1}^{(i,1:N)}$ corresponding to each parameter sample $\boldsymbol{\theta}_{t-1}^{(i)}$; that is, we need an ensemble of ensembles. Specifically, assume that the joint filtering distribution of parameters and states at time $t-1$ is given by

$$[\boldsymbol{\theta}_{t-1}, \mathbf{x}_{t-1}|\mathbf{z}_{1:t-1}] = \tfrac{1}{MN} \textstyle\sum_{i=1}^{M} \sum_{j=1}^{N} \delta_{(\boldsymbol{\theta}_{t-1}^{(i)}, \mathbf{x}_{t-1|t-1}^{(i,j)})}(\boldsymbol{\theta}_{t-1}, \mathbf{x}_{t-1}),$$

the marginal filtering distribution of $\boldsymbol{\theta}_{t-1}$ is given by $[\boldsymbol{\theta}_{t-1}|\mathbf{z}_{1:t-1}] = \frac{1}{M} \sum_{i=1}^{M} \delta_{(\boldsymbol{\theta}_{t-1}^{(i)})}(\boldsymbol{\theta}_{t-1})$, and the conditional filtering distribution of $\mathbf{x}_{t-1}$ given $\boldsymbol{\theta}_{t-1}$ is

$$[\mathbf{x}_{t-1}|\boldsymbol{\theta}_{t-1}^{(i)}, \mathbf{z}_{1:t-1}] = \tfrac{1}{N} \textstyle\sum_{j=1}^{N} \delta_{(\mathbf{x}_{t-1|t-1}^{(i,j)})}(\mathbf{x}_{t-1}).$$

The prior distribution at time t is then given by

$$[\boldsymbol{\theta}_t, \mathbf{x}_t|\boldsymbol{\theta}_{1:t-1}, \mathbf{z}_{1:t-1}] = [\boldsymbol{\theta}_t|\boldsymbol{\theta}_{1:t-1}, \mathbf{z}_{1:t-1}][\mathbf{x}_t|\boldsymbol{\theta}_{1:t}, \mathbf{z}_{1:t-1}] = [\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}][\mathbf{x}_t|\boldsymbol{\theta}_{1:t}, \mathbf{z}_{1:t-1}],$$

where $[\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}] = p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1})$ and

$$\begin{aligned}[\mathbf{x}_t|\boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}] &= \int [\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}] dP(\mathbf{x}_{t-1}|\boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}) \\ &= \tfrac{1}{N} \textstyle\sum_{j=1}^{N} \mathcal{N}_n(\mathbf{x}_t|\widetilde{\mathbf{x}}_{t|t-1}^{(i,j)}, \mathbf{Q}_t) \approx \mathcal{N}_n(\mathbf{x}_t|\widehat{\boldsymbol{\mu}}_{t|t-1}^{(i)}, \widehat{\boldsymbol{\Sigma}}_{t|t-1}^{(i)}), \end{aligned} \tag{13}$$

where we have used the same EnKF approximation as in (10). Note that $\widehat{\boldsymbol{\mu}}_{t|t-1}^{(i)}$ and $\widehat{\boldsymbol{\Sigma}}_{t|t-1}^{(i)}$ are obtained as described in Section 2.2 above using $\mathbf{Q}_t = \mathbf{Q}_t(\boldsymbol{\theta}_t)$, but only based on the ith forecast ensemble $\widetilde{\mathbf{x}}_{t|t-1}^{(i,1:N)}$, where $\widetilde{\mathbf{x}}_{t|t-1}^{(i,j)} = \mathcal{M}_t(\mathbf{x}_{t-1|t-1}^{(i,j)}; \boldsymbol{\theta}_t)$, $j = 1, \ldots, N$.

Based on the approximation in (13), the joint distribution of all variables at time t is given by

$$\begin{aligned}[\mathbf{z}_t, \mathbf{y}_t, \mathbf{x}_t, \boldsymbol{\theta}_t|\boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}] &= [\mathbf{z}_t|\mathbf{y}_t, \mathbf{x}_t, \boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}][\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}] \\ &\quad [\mathbf{x}_t|\boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}][\boldsymbol{\theta}_t|\boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}] \\ &= f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t) \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\mathbf{x}_t, \mathbf{R}_t) \mathcal{N}_n(\mathbf{x}_t|\widehat{\boldsymbol{\mu}}_{t|t-1}^{(i)}, \widehat{\boldsymbol{\Sigma}}_{t|t-1}^{(i)}) p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}^{(i)}), \end{aligned} \tag{14}$$

where $\mathbf{H}_t = \mathbf{H}_t(\boldsymbol{\theta}_t)$ and $\mathbf{R}_t = \mathbf{R}_t(\boldsymbol{\theta}_t)$ may also depend on the parameters $\boldsymbol{\theta}_t$.

3.1 Forecast independence

An important special case is given by forecast independence of states and parameters (Frei and Künsch, 2012): $[\mathbf{x}_t, \boldsymbol{\theta}_t|\mathbf{z}_{1:t-1}] = [\mathbf{x}_t|\mathbf{z}_{1:t-1}][\boldsymbol{\theta}_t|\mathbf{z}_{1:t-1}]$. Forecast independence leads to enormous

simplifications in the algorithms described later, because the prior distribution of the states can be calculated based on the entire ensemble of ensembles from the previous filtering distributions, instead of using only the corresponding ensemble as in (13); that is,

$$[\mathbf{x}_t|\boldsymbol{\theta}_t, \mathbf{z}_{1:t-1}] = [\mathbf{x}_t|\mathbf{z}_{1:t-1}] = \int [\mathbf{x}_t|\mathbf{x}_{t-1}] dP(\mathbf{x}_{t-1}|\mathbf{z}_{1:t-1}) = \mathcal{N}_n(\mathbf{x}_t|\widehat{\boldsymbol{\mu}}_{t|t-1}, \widehat{\boldsymbol{\Sigma}}_{t|t-1}),$$

where $\widehat{\boldsymbol{\mu}}_{t|t-1}$ and $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$ are now based on all ensemble members $\widetilde{\mathbf{x}}_{t|t-1}^{(1:M,1:N)}$. Because this means that we do not need a separate ensemble for each previous parameter value, we can set $N = 1$ and write the forecast ensemble as $\widetilde{\mathbf{x}}_{t|t-1}^{(1:M)}$. Also, because the prior distribution does not depend on $\boldsymbol{\theta}_t$, we only have to propagate the ensemble once from time $t-1$ to t in the algorithms described later, which is especially important if $\mathcal{M}_t$ is computationally expensive.

Forecast independence can be proved asymptotically if the parameters $\boldsymbol{\theta}_t$ only appear in $\mathbf{R}_t$ (see Frei and Künsch, 2012). It also arises (exactly) if the parameters do not appear in the evolution layer (3) and are temporally indendent (i.e., $p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}) = p_t(\boldsymbol{\theta}_t)$), because then

$$[\mathbf{x}_t|\boldsymbol{\theta}_t, \mathbf{z}_{1:t-1}] = \int [\mathbf{x}_t|\mathbf{x}_{t-1}, \boldsymbol{\theta}_t] dP(\mathbf{x}_{t-1}|\boldsymbol{\theta}_t, \mathbf{z}_{1:t-1}) = [\mathbf{x}_t|\mathbf{z}_{1:t-1}],$$

as above.

3.2 Approximation of the likelihood for known $\mathbf{y}_t$

It is also possible to obtain an approximation of the likelihood based on the EnKF. We first study likelihood approximations for the special case of known $\mathbf{y}_t$ (i.e., the hierarchical SSM given by (2)–(4)).

The distribution of all data up to time t is then given by,

$$[\mathbf{y}_{1:t}|\boldsymbol{\theta}_{1:t}] = \prod_{k=1}^{t} [\mathbf{y}_k|\boldsymbol{\theta}_{1:t}, \mathbf{y}_{1:k-1}] = \prod_{k=1}^{t} [\mathbf{y}_k|\boldsymbol{\theta}_{1:k}, \mathbf{y}_{1:k-1}],$$

where we have now made the conditioning on the parameters explicit. However, the in high-dimensional setting with potentially nonlinear evolution considered here, the filtering likelihood at time t, $\mathcal{L}_t(\boldsymbol{\theta}_{1:t}) := [\mathbf{y}_t|\boldsymbol{\theta}_{1:t}, \mathbf{y}_{1:t-1}]$, is not available in closed form and cannot be exactly evaluated. Instead, in any parameter-inference procedure, we need to estimate this likelihood for given fixed parameter values based on the forecast ensemble $\widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}$. For the remainder of Section 3.2, we focus on a single time point t and assume that $\widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}$ is a sample from the true forecast distribution, $[\mathcal{M}_t(\mathbf{x}_{t-1})|\mathbf{y}_{1:t-1}, \boldsymbol{\theta}_{1:t}]$.

For Bayesian sampling-based inference, the closer the likelihood estimator is to the true likelihood, the better the performance of the parameter-inference procedure. Specifically, unbiased estimators of the true likelihood lead to so-called exact approximate procedures that can be shown to sample from the correct target distribution; but the variability of the estimator is also crucial, in that the variance of the logarithm of the estimator of the likelihood determines the performance of the parameter inference for finite computation time (Andrieu and Roberts, 2009). Variances smaller than 2 are often recommended for satisfactory performance (e.g., Doucet et al., 2015).

3.2.1 Particle likelihood

If the prior distribution in (7) were used, the filtering likelihood at time t would be a mixture of normals,

$$[\mathbf{y}_t|\boldsymbol{\theta}_{1:t}, \mathbf{y}_{1:t-1}] = \tfrac{1}{N}\sum_{j=1}^{N} \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\widetilde{\mathbf{x}}_{t|t-1}^{(j)}, \mathbf{H}_t\mathbf{Q}_t\mathbf{H}_t' + \mathbf{R}_t) =: \mathcal{L}_t^{\mathrm{P}}(\mathbf{y}_t|\boldsymbol{\theta}_t, \widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}). \tag{15}$$

Except for the fact that we are able to integrate out the innovation error with covariance $\mathbf{Q}_t$ here due to the linear Gaussian observations, $\mathcal{L}_t^{\mathrm{P}}$ is the likelihood approximation used in several inference procedures based on particle filters (e.g., Andrieu et al., 2010; Johansen et al., 2012) and in some approaches in the EnKF literature (e.g., Ueno and Nakamura, 2014, 2016). Using the notation from (9), we can write this likelihood as $\mathcal{L}_t^{\mathrm{P}}(\mathbf{y}_t|\boldsymbol{\theta}_t, \widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}) \propto (1/N)\sum_{j=1}^{N} \alpha_t^{(j)}$. While the $\alpha_t^{(j)}$ and thus also the particle likelihood $\mathcal{L}_t^{\mathrm{P}}$ are unbiased for the true likelihood (Del Moral, 2004), leading to exact approximate algorithms, the variance of the particle likelihood can be very large in high dimensions, resulting in algorithms with potentially poor performance in finite computation time. We now illustrate this using a simple example.

Example 1. *For a single time point, consider the forecast distribution* $\widetilde{\mathbf{x}} \sim \mathcal{N}_n(\mathbf{0}, \kappa\mathbf{I}_n)$*, with model error covariance* $\mathbf{Q} = \mathbf{0}$*, noise covariance* $\mathbf{R} = \theta\mathbf{I}_n$*, observation matrix* $\mathbf{H} = \mathbf{I}_n$*, and thus* $n = m$*; that is,* $\mathbf{y} \sim \mathcal{N}_n(\mathbf{0}, (\kappa + \theta)\mathbf{I}_n)$*. The unknown parameter is* θ*.*

In the setting of Example 1, Figure 2 shows the distribution of the particle likelihood $\mathcal{L}_t^{\mathrm{P}}$, which becomes increasingly skewed for increasing n. Further, it shows the trace plot for a "pseudo-marginal" Metropolis-Hastings (MH) sampler targeting θ for $n = 50$ using $\mathcal{L}_t^{\mathrm{P}}$, whose large variances cause the Markov chain to get "stuck" after a particularly large likelihood is obtained. Finally, we explore the ensemble size N necessary to keep the variance of the particle loglikelihood (averaged over 100 realizations of $\mathbf{y}$) below 2, and this number clearly scales exponentially with the dimension n. This result can also be shown analytically:

Proposition 1. *In the setting of Example 1, we have* $\mathrm{var}(\log \mathcal{L}_t^{\mathrm{P}}(\mathbf{y}|\theta, \widetilde{\mathbf{x}}^{(1:N)})) = \mathcal{O}(e^n/N)$ *as* $n, N \to \infty$*, where the variance is evaluated with respect to* $\widetilde{\mathbf{x}}^{(1:N)}$*, and the* $\widetilde{\mathbf{x}}^{(j)}$ *are iid from the true forecast distribution.*

All proofs are provided in Appendix A.

The results obtained here for a particle or importance-sampling likelihood at a single time point extend directly to a particle filter in a state-space setting. While a particle filter used to approximate a likelihood in a sequential setting will typically include a resampling step at every time point (e.g., Andrieu et al., 2010), the resulting forecast or prior distribution at the next time point can be at most as good as the iid sampling from the true prior considered here (assuming that $\mathcal{M}_t$ is intractable), and so the particle-filter approximation to the likelihood will also degenerate in high dimensions.

3.2.2 EnKF likelihood

Due to the problems of the PF likelihood in high dimensions, we use the EnKF approximation of the prior distribution in (10) to obtain an EnKF likelihood:

$$\begin{aligned}
[\mathbf{y}_t|\boldsymbol{\theta}_{1:t}, \mathbf{y}_{1:t-1}] &= \textstyle\int [\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}_{1:t}, \mathbf{y}_{1:t-1}] dP(\mathbf{x}_t|\boldsymbol{\theta}_{1:t}, \mathbf{y}_{1:t-1}) \\
&= \textstyle\int \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\mathbf{x}_t, \mathbf{R}_t)\mathcal{N}_n(\mathbf{x}_t|\widehat{\boldsymbol{\mu}}_{t|t-1}, \widehat{\boldsymbol{\Sigma}}_{t|t-1})d\mathbf{x}_t \\
&= \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\widehat{\boldsymbol{\mu}}_{t|t-1}, \mathbf{H}_t\widehat{\boldsymbol{\Sigma}}_{t|t-1}\mathbf{H}_t' + \mathbf{R}_t) =: \mathcal{L}_t^{\mathrm{E}}(\mathbf{y}_t|\boldsymbol{\theta}_t, \widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}),
\end{aligned} \tag{16}$$

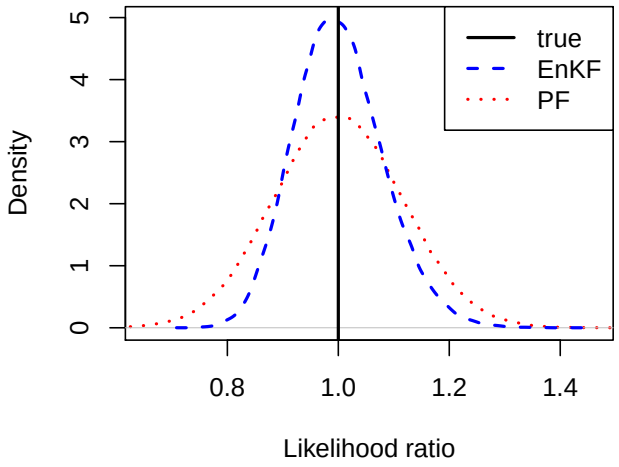

(a) Distribution of likelihoods for $n = 1$

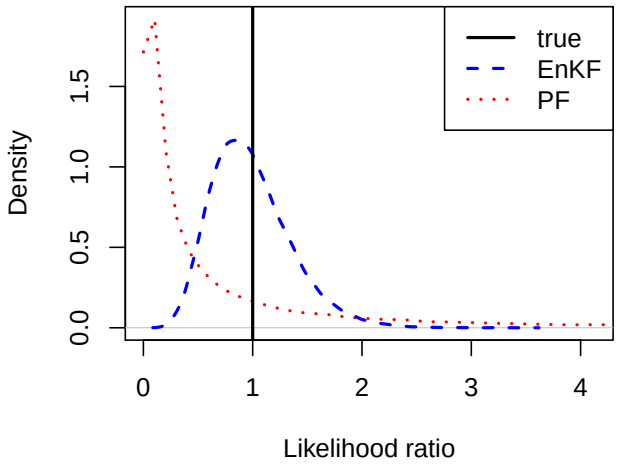

(b) Distribution of likelihoods for $n = 6$

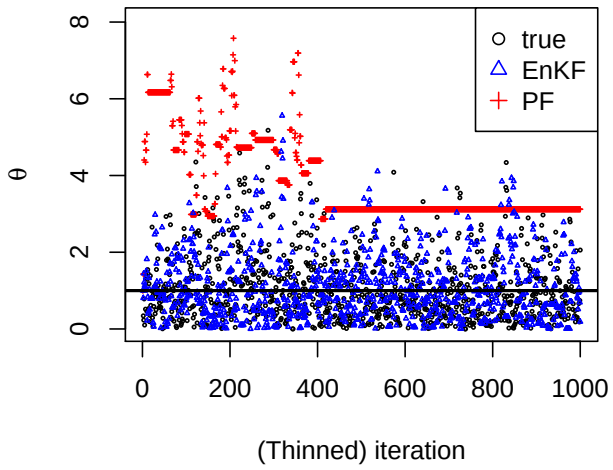

(c) Trace plots for MH (thinned by factor 5)

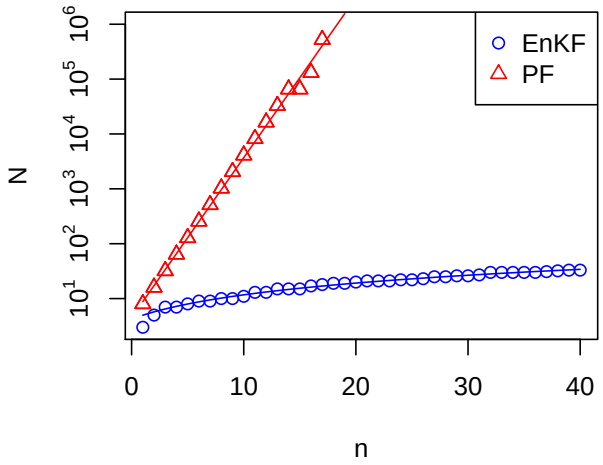

(d) Necessary N to keep the variance of the loglikelihood below 2

Figure 2: Comparison of the likelihoods discussed in Section 3.2 in the setting of Example 1 with $\kappa = 4$, and true parameter value $\theta = 1$. Panels (a) and (b): Distributions of EnKF and PF likelihood approximations (divided by the true likelihood) at $\theta = 1$ for $N = 50$ with fixed data $\mathbf{y}$ and varying forecast ensembles. Panel (c): Trace plot for "pseudo-marginal" Metropolis-Hastings (MH) algorithms targeting θ with a uniform prior on $\mathbb{R}^+$, $n = N = 50$, and a normal proposal distribution with standard deviation 0.8. Panel (d): Minimum N to keep the variance of the loglikelihood below 2, along with a least-squares line fitted on the original scale (for EnKF) and on a log-scale (for PF).

where $\mathbf{H}_t$, $\mathbf{R}_t$, and $\widehat{\boldsymbol{\mu}}_{t|t-1}$ and $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$ (through $\mathcal{M}_t$ and $\mathbf{Q}_t$) may depend on $\boldsymbol{\theta}_t$.

The EnKF likelihood has been successfully used in several examples and applications (e.g., Mitchell and Houtekamer, 2000; Stroud et al., 2010; Frei and Künsch, 2012; Stroud et al., 2016). For linear SSMs, the true forecast distribution and likelihood are both exactly Gaussian, and so the EnKF likelihood converges to the true likelihood as $N \to \infty$ (assuming no localization for $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$). For finite N, in contrast to the PF likelihood, the EnKF likelihood is a biased estimator of the true likelihood, and thus does not lead to exact approximate inference. Special care should be taken regarding parameters influenced by tapering, as, for example, a spatial-scale parameter in the innovation covariance $\mathbf{Q}_t$ might be overestimated to compensate for an overly restrictive taper in the forecast covariance $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$ (see Section 2.2). The estimated parameter is then useful for running an EnKF with the same taper settings, but it might not be the "true" spatial scale of the innovation covariance in the model (1)–(4). However, more importantly, the variance of the EnKF (log)likelihood can be relatively small even in high dimensions, as long as a good (regularized) estimator $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$ can be found. This can be viewed as a manifestation of the well-known trade-off between bias and variance that is central to statistical learning with "big" data (e.g., Hastie et al., 2009).

To illustrate this last point, we consider the EnKF likelihood in the setting of Example 1 with a diagonal $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$. Figure 2 shows that the EnKF likelihood is less skewed, less variable, and thus closer to the true likelihood than the particle likelihood. This also leads to much better mixing for the MH sampler, and the resulting estimated posterior for θ (not shown) is nearly indistinguishable from the estimated posterior using the true likelihood. Figure 2(d) indicates that for the EnKF likelihood, N only has to increase linearly with the dimension n to keep the variance of the loglikelihood bounded. This can also be shown analytically:

Proposition 2. *In the setting of Example 1, we have* $\operatorname{var}(\log \mathcal{L}_t^{\mathrm{E}}(\mathbf{y}|\boldsymbol{\theta}, \widetilde{\mathbf{x}}^{(1:N)})) = \mathcal{O}(n/N)$ *as* $n, N \to \infty$*, where the variance is evaluated with respect to* $\widetilde{\mathbf{x}}^{(1:N)}$*, and the* $\widetilde{\mathbf{x}}^{(j)}$ *are iid from the true forecast distribution.*

To explore whether this result holds more broadly than just in the independent case, we repeated the simulations in Figure 2(d) for a slightly modified version of Example 1 with a tridiagonal forecast covariance matrix. This produced a plot very similar to Figure 2(d).

If the true forecast distribution (and thus also the true likelihood) is non-Gaussian, the EnKF likelihood converges (for $N \to \infty$) to the Dawid–Sebastiani scoring rule (Dawid and Sebastiani, 1999), and so, from a frequentist perspective, parameter inference based on the EnKF likelihood at a single time point can be interpreted as (approximate) minimum contrast estimation (Birgé and Massart, 1993; Pfanzagl, 1969) under the Dawid–Sebastiani score. Due to the the propriety of this scoring rule, the minimum-contrast estimator can be shown, using the proof of Wald (1949), to be consistent under regularity conditions (cf. Gneiting and Katzfuss, 2014, Sect. 3.4.2). From a Bayesian perspective, inference based on the EnKF likelihood essentially amounts to computing the (approximate) "Gibbs posterior" using the Dawid–Sebastiani as the empirical risk function (Jiang and Tanner, 2008; Li et al., 2013; Katzfuss and Bhattacharya, 2015). Under certain conditions, most importantly that $\hat{\boldsymbol{\theta}}_t := \arg\max \mathcal{L}_t^{\mathrm{E}}(\mathbf{y}_t|\boldsymbol{\theta}_t, \widetilde{\mathbf{x}}_t^{(1:N)})$ is consistent and asymptotically normal as $m_t \to \infty$, it might be possible to scale and rotate posterior samples of $\boldsymbol{\theta}_t$ to reflect the true asymptotic uncertainty in the frequentist sense (Shaby, 2014), but this will be pursued in future work.

3.3 Integrated likelihood

We now extend the EnKF likelihood for known $\mathbf{y}_t$ from Section 3.2.2 to the full HSSM with the general transformation layer (1): $\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t \sim f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t)$. This integrated EnKF likelihood of the data $\mathbf{z}_t$ at time t, with $\mathbf{x}_t$ and $\mathbf{y}_t$ integrated out, is an important component of some of the algorithms described later. Using (14), we have

$$\begin{aligned}[\mathbf{z}_t|\boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}] &= \int f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t) \int \mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t\mathbf{x}_t, \mathbf{R}_t)[\mathbf{x}_t|\boldsymbol{\theta}_t, \boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}]d\mathbf{x}_t d\mathbf{y}_t \\ &= \int f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t)\mathcal{L}_t^{\mathrm{E}}(\mathbf{y}_t|\boldsymbol{\theta}_t, \widetilde{\mathbf{x}}_{t|t-1}^{(i,1:N)})d\mathbf{y}_t =: \mathcal{L}_t^{\mathrm{Z}}(\mathbf{z}_t|\boldsymbol{\theta}_t, \widetilde{\mathbf{x}}_{t|t-1}^{(i,1:N)}), \end{aligned} \tag{17}$$

where $\mathcal{L}_t^{\mathrm{E}}$ is defined in (16), and $\mathbf{H}_t$, $\mathbf{R}_t$, and $\widehat{\boldsymbol{\mu}}_{t|t-1}^{(i)}$ and $\widehat{\boldsymbol{\Sigma}}_{t|t-1}^{(i)}$ may depend on $\boldsymbol{\theta}_t$.

Evaluating this likelihood requires carrying out an integral over the m_t-dimensional vector $\mathbf{y}_t$. Strategies for doing so depend on the choice of transformation distribution f_t. If the data $\mathbf{z}_t$ are Gaussian and the transformation is simply the identity, the integral degenerates and analytical integration is trivial. If it is possible to sample from $[\mathbf{y}_t|\mathbf{z}_t, \mathbf{x}_t, \boldsymbol{\theta}_t]$, the integral can be approximated using these posterior samples (e.g., Chib and Jeliazkov, 2001; Raftery et al., 2007). Finally, if $f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t)$ is a continous function in $\mathbf{y}_t$, a Laplace approximation (Tierney and Kadane, 1986) can be used to approximate the integral. As we are already approximating $\mathbf{y}_t$ as a latent Gaussian field in $\mathcal{L}_t^{\mathrm{E}}(\mathbf{y}_t|\boldsymbol{\theta}_t, \widetilde{\mathbf{x}}_{t|t-1}^{(i,1:N)})$, the Laplace approximation of the integral should often result in negligible additional error, especially if the observations in $\mathbf{z}_t$ are independent given $\mathbf{y}_t$ and $\boldsymbol{\theta}_t$ (e.g., Rue et al., 2009); see Section 5.2 for an example.

4 Extended ensemble Kalman filters and smoothers for hierarchical state-space models

For HSSMs of the form (1)–(4) with non-Gaussian observations or unknown parameters, we now propose extended EnKF methods. The general idea is to approximate the inference on the state vector in suitable existing Bayesian inference techniques using the EnKF or the EnKS.

4.1 Extended ensemble Kalman filters

We begin by discussing extended EnKFs that are suitable for filtering inference, which at each time t aims to find the filtering distribution $[\mathbf{x}_t, \boldsymbol{\theta}_t|\mathbf{z}_{1:t}]$ given the data $\mathbf{z}_{1:t}$ collected up to time t. Smoothing inference is discussed in Section 4.2 below.

Throughout this section, we will illustrate our algorithms using the following simple toy example, which is the same as Example 1 except with a general state covariance matrix $\boldsymbol{\Sigma}$:

Example 2. *At a single time point, consider the model* $\mathbf{y}|\mathbf{x}, \theta \sim \mathcal{N}_n(\mathbf{x}, \theta\mathbf{I})$, $\mathbf{x} \sim \mathcal{N}_n(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, $\theta \sim \mathcal{IG}(a, b)$, *and we observe* $\mathbf{y}$ *directly (i.e.,* $\mathbf{z} = \mathbf{y}$*). Further, we assume forecast independence, so that* $[\mathbf{x}, \theta] = [\mathbf{x}][\theta]$. *Thus, the joint distribution from* (14) *is* $[\mathbf{y}, \mathbf{x}, \theta] = \mathcal{N}_n(\mathbf{y}|\mathbf{x}, \theta\mathbf{I})\mathcal{N}_n(\mathbf{x}|\boldsymbol{\mu}, \boldsymbol{\Sigma})\mathcal{IG}(\theta|a, b)$. *The goal is to sample from the joint posterior distribution,* $[\mathbf{x}, \theta|\mathbf{y}]$.

4.1.1 Gibbs ensemble Kalman filter (GEnKF)

First, we consider a combination of an EnKF and a Gibbs sampler (Geman and Geman, 1984; Gelfand and Smith, 1990) for filtering inference in HSSMs. For this algorithm we assume forecast independence as described in Section 3.1, which enables inference based on a single ensemble $\mathbf{x}_{t|t}^{(1:M)} = \{\mathbf{x}_{t|t}^{(1)}, \ldots, \mathbf{x}_{t|t}^{(M)}\}$ (i.e., $N = 1$). At each time t, our proposed Gibbs ensemble Kalman filter (GEnKF) iteratively samples $\mathbf{y}_t$, $\mathbf{x}_t$, and $\boldsymbol{\theta}_t$ from their respective full conditional distributions (FCDs), which are all proportional to the joint distribution in (14):

Algorithm 2: Gibbs ensemble Kalman filter (GEnKF)
Starting with an initial ensemble $\{(\boldsymbol{\theta}_0^{(i)}, \mathbf{x}_{0|0}^{(i)}) : i = 1, \ldots, M\}$, where $(\boldsymbol{\theta}_0^{(i)}, \mathbf{x}_{0|0}^{(i)}) \sim [\boldsymbol{\theta}_0, \mathbf{x}_0] = p_0(\boldsymbol{\theta}_0)\mathcal{N}_n(\mathbf{x}_0|\boldsymbol{\mu}_{0|0}, \boldsymbol{\Sigma}_{0|0})$, do the following for each $t = 1, 2, \ldots$:

1. Obtain the forecast ensemble as $\widetilde{\mathbf{x}}_{t|t-1}^{(i)} = \mathcal{M}_t(\mathbf{x}_{t-1|t-1}^{(i)})$, for all $i = 1, \ldots, M$.
2. Find starting values for $\mathbf{y}_t^{(i)}$ and $\boldsymbol{\theta}_t^{(i)}$ for all $i = 1, \ldots, M$.
3. For $i = 1, \ldots, M$, iterate between the following steps until convergence:
   (a) Sample $\mathbf{x}_{t|t}^{(i)}$ from $[\mathbf{x}_t|\mathbf{y}_{1:t}^{(i)}, \boldsymbol{\theta}_t^{(i)}]$ by computing the EnKF update $\mathbf{x}_{t|t}^{(i)} = \mathcal{U}_t(\widetilde{\mathbf{x}}_{t|t-1}^{(i)}|\mathbf{y}_t^{(i)}, \widetilde{\mathbf{x}}_{t|t-1}^{(1:M)}, \boldsymbol{\theta}_t^{(i)})$ (see (12)).
   (b) Sample $\mathbf{y}_t^{(i)}$ from $[\mathbf{y}_t|\mathbf{x}_{t|t}^{(i)}, \boldsymbol{\theta}_t^{(i)}, \mathbf{z}_{1:t}] \propto f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t^{(i)})\mathcal{N}_{m_t}(\mathbf{y}_t|\mathbf{H}_t(\boldsymbol{\theta}_t^{(i)})\mathbf{x}_{t|t}^{(i)}, \mathbf{R}_t(\boldsymbol{\theta}_t^{(i)}))$.
   (c) Sample $\boldsymbol{\theta}_t^{(i)}$ from $[\boldsymbol{\theta}_t|\mathbf{y}_t^{(i)}, \mathbf{x}_{t|t}^{(i)}, \mathbf{z}_{1:t}] \propto f_t(\mathbf{z}_t|\mathbf{y}_t^{(i)}, \boldsymbol{\theta}_t)\mathcal{N}_{m_t}(\mathbf{y}_t^{(i)}|\mathbf{H}_t(\boldsymbol{\theta}_t)\mathbf{x}_{t|t}^{(i)}, \mathbf{R}_t(\boldsymbol{\theta}_t))p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}^{(i)})$.

   After convergence, $(\mathbf{x}_{t|t}^{(i)}, \boldsymbol{\theta}_t^{(i)})$ is a joint sample from $[\mathbf{x}_t, \boldsymbol{\theta}_t|\mathbf{z}_{1:t}]$, and $\mathbf{x}_{t|t}^{(1:M)}$ is the filtering ensemble at time t.

Example 3. *For the simple model in Example 2, the GEnKF proceeds as follows:*

1. *Generate $\widetilde{\mathbf{x}}^{(i)} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, for all $i = 1, \ldots, M$, and compute $\widetilde{\boldsymbol{\Sigma}}$, the regularized sample covariance of $\widetilde{\mathbf{x}}^{(1:M)}$.*
2. *Find starting values for $\theta^{(i)}$ for all $i = 1, \ldots, M$.*
3. *For $i = 1, \ldots, M$, iterate between the following steps until convergence:*
   (a) Set $\mathbf{x}^{(i)} = \widetilde{\mathbf{x}}^{(i)} + \widetilde{\boldsymbol{\Sigma}}(\widetilde{\boldsymbol{\Sigma}} + \theta^{(i)}\mathbf{I})^{-1}(\mathbf{y} - \widetilde{\mathbf{x}}^{(i)} - \mathbf{v}^{(i)})$, where $\mathbf{v}^{(i)} \sim \mathcal{N}(\mathbf{0}, \theta^{(i)}\mathbf{I})$.
   (b) Sample $\theta^{(i)} \sim \mathcal{IG}(a + \frac{n}{2}, b + \frac{1}{2}(\mathbf{y} - \mathbf{x}^{(i)})'(\mathbf{y} - \mathbf{x}^{(i)}))$.

   After convergence, $(\mathbf{x}^{(i)}, \theta^{(i)})$ is a sample from $[\mathbf{x}, \theta|\mathbf{y}]$.

Thus, at a particular time point, the GEnKF update is similar to a standard Gibbs sampler that iterates between $[\mathbf{x}|\theta, \mathbf{y}]$ and $[\theta|\mathbf{x}, \mathbf{y}]$. The only difference is that instead of sampling directly from $[\mathbf{x}|\theta, \mathbf{y}]$, our algorithm samples approximately from $[\mathbf{x}|\theta, \mathbf{y}]$ using the EnKF update. Thus, in contrast to the Gibbs sampler, the GEnKF can be applied to high-dimensional and nonlinear SSMs.

Note that iterating through Step 3(a)–(c) of Algorithm 2 can be done completely independently (i.e., in parallel) for $i = 1, \ldots, M$. As the GEnKF requires sampling from the FCDs of $\mathbf{y}_t$ and $\boldsymbol{\theta}_t$ in (b) and (c), respectively, the algorithm is most efficient if the distributions $f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t)$ and $p_t(\boldsymbol{\theta}_t)$ are such that these FCDs are available in closed form. Otherwise, Metropolis updates are necessary for these variables, which is generally only feasible if the number of parameters in $\boldsymbol{\theta}_t$ is small, and if the data in $\mathbf{z}_t$ are independent given the state $\mathbf{x}_t$. Several examples are presented in Section 5.1.

For computational feasibility of Algorithm 2, it is crucial to find good starting values for $\mathbf{y}_t$ and $\boldsymbol{\theta}_t$, in order to minimize the required number of iterations. In some applications, it is natural to use the values from the previous time point (i.e., set the initial values as $\mathbf{y}_t^{(i)} = \mathbf{y}_{t-1}^{(i)}$ and $\boldsymbol{\theta}_t^{(i)} = \boldsymbol{\theta}_{t-1}^{(i)}$). Another possibility is to choose the value of $\boldsymbol{\theta}_t$ that maximizes the integrated likelihood in (17). Note that in machine learning, often only a few or even a single update (called contrastive divergence) is used to sample from restricted Boltzman machine Markov random fields for training deep neural networks (Hinton, 2002). Although this adds bias, empirical and some theoretical support suggests that it is an effective and efficient approach in the deep learning context (e.g., Bengio and Delalleau, 2009).

If $\boldsymbol{\theta}_t$ is a static parameter (i.e., $\boldsymbol{\theta} = \boldsymbol{\theta}_1 = \boldsymbol{\theta}_2 = \ldots$), Algorithm 2 will fail, in that the distribution in Step 3(c) will be a point mass at the previous parameter value, but this can be remedied by introducing artificial innovation noise (e.g., Kitagawa, 1998; Liu and West, 2001) to make Algorithm 2 feasible.

Finally, a Gibbs ensemble Kalman filter could in principle also be applied if forecast independence does not hold. This is difficult, however, because $[\boldsymbol{\theta}_t, \mathbf{x}_t|\mathbf{z}_{1:t}] = \int [\boldsymbol{\theta}_t, \mathbf{x}_t|\boldsymbol{\theta}_{t-1}, \mathbf{z}_{1:t}] dP(\boldsymbol{\theta}_{t-1}|\mathbf{z}_{1:t})$, where it is straightforward to sample from $[\boldsymbol{\theta}_t, \mathbf{x}_t|\boldsymbol{\theta}_{t-1}, \mathbf{z}_{1:t}]$ using a Gibbs sampler as above, but it is highly challenging to compute the probability or weight $P(\boldsymbol{\theta}_{t-1} = \boldsymbol{\theta}_{t-1}^{(i)}|\mathbf{z}_{1:t})$, which will also degenerate in high dimensions. Thus we will not pursue this general case here.

Perhaps the approaches in the literature most closely related to our GEnKF are Myrseth and Omre (2010) and Tsyrulnikov and Rakitko (2015), who model the forecast covariance matrix as random with an inverse-Wishart prior. This matrix is not a model parameter, which makes Bayesian inference difficult. Myrseth and Omre (2010) propose a Gibbs-like algorithm with a single iteration, but they do not take the data into account when calculating the posterior of the forecast covariance. Tsyrulnikov and Rakitko (2015) seem to ignore the fact that forecast independence does not hold in their model.

4.1.2 Particle ensemble Kalman filter (PEnKF)

For high-dimensional SSMs, particle filters typically degenerate (e.g., Doucet et al., 2001; Snyder et al., 2008). While Rao-Blackwellized particle filters (RBPFs; Doucet et al., 2000a) can scale to higher dimensions by integrating out part of the state analytically, for example using a Kalman filter (Chen and Liu, 2000), the integral is generally not available for HSSMs with nonlinear evolution or non-Gaussian data. An exact approximation of RBPFs proposed in Johansen et al. (2012) uses "local" particle filters to carry out the integration, thus employing the particle likelihood from Section 3.2.1, which was shown to scale poorly as well.

As an alternative, we propose here a particle ensemble Kalman filter (PEnKF), which uses the EnKF likelihood from Sections 3.2.2 and 3.3 to (approximately) integrate out the high-dimensional state. The PEnKF is suitable for HSSMs for which forecast independence does not hold, or for which efficient sampling from the FCDs in Step 3 of Algorithm 2 is not possible. As

a particle filter is applied to $\boldsymbol{\theta}_t$, the PEnKF will work best if the number of unknown parameters is not too high. This algorithm does not require forecast independence for the parameters as the GEnKF above, although forecast independence does result in considerable simplifications (see below).

Algorithm 3: Particle ensemble Kalman filter (PEnKF)
Initialize the algorithm with an equally weighted ensemble $(\boldsymbol{\theta}_0^{(i)}, \mathbf{x}_{0|0}^{(i,j)}) \sim [\boldsymbol{\theta}_0, \mathbf{x}_0] = p_0(\boldsymbol{\theta}_0)\mathcal{N}_n(\mathbf{x}_0|\boldsymbol{\mu}_{0|0}, \boldsymbol{\Sigma}_{0|0})$ with $w_0^{(i)} = 1/M$, $i = 1, \ldots, M; j = 1, \ldots, N$. Then, for $t = 1, 2, \ldots$:

1. For $i = 1, \ldots, M$:
   (a) Sample a particle $\boldsymbol{\theta}_t^{(i)}$ from a proposal distribution $q_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}^{(i)})$.
   (b) Propagate the corresponding ensemble: $\widetilde{\mathbf{x}}_{t|t-1}^{(i,j)} = \mathcal{M}_t(\mathbf{x}_{t-1|t-1}^{(i,j)}; \boldsymbol{\theta}_t^{(i)})$, $j = 1, \ldots, N$.
   (c) Calculate the particle weight:
   $w_t^{(i)} \propto w_{t-1}^{(i)} \mathcal{L}_t^{Z}(\mathbf{z}_t|\boldsymbol{\theta}_t^{(i)}, \widetilde{\mathbf{x}}_{t|t-1}^{(i,1:N)}) p_t(\boldsymbol{\theta}_t^{(i)}|\boldsymbol{\theta}_{t-1}^{(i)}) / q_t(\boldsymbol{\theta}_t^{(i)}|\boldsymbol{\theta}_{t-1}^{(i)})$.
   (d) Generate $\mathbf{x}_{t|t}^{(i,j)}$ from $[\mathbf{x}_t|\mathbf{z}_{1:t}, \boldsymbol{\theta}_{1:t}^{(i)}]$ for $j = 1, \ldots, N$ (see below).
2. The filtering distribution is given by
   $[\boldsymbol{\theta}_t, \mathbf{x}_t|\mathbf{z}_{1:t}] \approx \sum_{i=1}^{M} w_t^{(i)} \frac{1}{N} \sum_{j=1}^{N} \delta_{(\boldsymbol{\theta}_t^{(i)}, \mathbf{x}_{t|t}^{(i,j)})}(\boldsymbol{\theta}_t, \mathbf{x}_t)$.
3. If desired, resample $(\boldsymbol{\theta}_t^{(i)}, \mathbf{x}_{t|t}^{(i,1:N)})$ to obtain an equally weighted sample (see, e.g., Douc et al., 2005, for a comparison of resampling schemes).

By noting that $[\boldsymbol{\theta}_{1:t}, \mathbf{z}_{1:t}] = [\boldsymbol{\theta}_t, \mathbf{z}_t|\boldsymbol{\theta}_{1:t-1}, \mathbf{z}_{1:t-1}][\boldsymbol{\theta}_{1:t-1}, \mathbf{z}_{1:t-1}]$, the expression for the weights in Step 1(c) is derived as follows:

$$w_t^{(i)} \propto \frac{w_{t-1}^{(i)}}{q_t(\boldsymbol{\theta}_t^{(i)}|\boldsymbol{\theta}_{t-1}^{(i)})} \frac{[\boldsymbol{\theta}_{1:t}^{(i)}, \mathbf{z}_{1:t}]}{[\boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}]} = \frac{w_{t-1}^{(i)}}{q_t(\boldsymbol{\theta}_t^{(i)}|\boldsymbol{\theta}_{t-1}^{(i)})} [\boldsymbol{\theta}_t^{(i)}, \mathbf{z}_t|\boldsymbol{\theta}_{1:t-1}^{(i)}, \mathbf{z}_{1:t-1}],$$

where $[\boldsymbol{\theta}_t, \mathbf{z}_t|\boldsymbol{\theta}_{1:t-1}, \mathbf{z}_{1:t-1}] = [\boldsymbol{\theta}_t|\boldsymbol{\theta}_{1:t-1}, \mathbf{z}_{1:t-1}][\mathbf{z}_t|\boldsymbol{\theta}_{1:t}, \mathbf{z}_{1:t-1}] = p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1})[\mathbf{z}_t|\boldsymbol{\theta}_{1:t}, \mathbf{z}_{1:t-1}]$, and $[\mathbf{z}_t|\boldsymbol{\theta}_{1:t}^{(i)}, \mathbf{z}_{1:t-1}]$ is the integrated likelihood in (17).

Example 4. *For the simple model in Example 2 (and ignoring forecast independence), the PEnKF first initializes weights $w_0^{(i)} = 1/M$ for $i = 1, \ldots, M$, and then proceeds as follows:*

1. *For $i = 1, \ldots, M$:*
   (a) Sample $\theta^{(i)}$ from a proposal distribution, $q(\theta)$.
   (b) Generate the corresponding ensemble $\widetilde{\mathbf{x}}^{(i,j)} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, $j = 1, \ldots, N$, and compute the sample mean $\widetilde{\boldsymbol{\mu}}^{(i)}$ and the (regularized) covariance matrix $\widetilde{\boldsymbol{\Sigma}}^{(i)}$ of $\widetilde{\mathbf{x}}^{(i,1:N)}$.
   (c) Calculate the particle weight: $w^{(i)} \propto w_0^{(i)} \mathcal{N}(\mathbf{y}|\widetilde{\boldsymbol{\mu}}^{(i)}, \widetilde{\boldsymbol{\Sigma}}^{(i)} + \theta^{(i)}\mathbf{I}) \mathcal{IG}(\theta^{(i)}|a, b) / q(\theta^{(i)})$.

(d) Set $\mathbf{x}^{(i,j)} = \widetilde{\mathbf{x}}^{(i,j)} + \widetilde{\boldsymbol{\Sigma}}^{(i)}(\widetilde{\boldsymbol{\Sigma}}^{(i)} + \theta^{(i)}\mathbf{I})^{-1}(\mathbf{y} - \widetilde{\mathbf{x}}^{(i,j)} - \mathbf{v}^{(i,j)})$*, where* $\mathbf{v}^{(i,j)} \sim \mathcal{N}(\mathbf{0}, \theta^{(i)}\mathbf{I})$*, for* $j = 1, \ldots, N$*.*

2. *The posterior distribution is given by* $[\theta, \mathbf{x}|\mathbf{y}] \approx \sum_{i=1}^{M} w^{(i)} \frac{1}{N} \sum_{j=1}^{N} \delta_{(\theta^{(i)}, \mathbf{x}^{(i,j)})}(\theta, \mathbf{x})$*.*

3. *If desired, resample* $(\theta^{(i)}, \mathbf{x}^{(i,1:N)})$ *with weights* $w^{(i)}$ *to obtain an equally-weighted sample.*

Thus, the PEnKF is similar to a Rao-Blackwellized particle-filter, except that the high-dimensional state $\mathbf{x}$ *is integrated out and sampled using the EnKF approximation. This again allows application to high-dimensional (and nonlinear) SSMs.*

For Step 2 in Algorithm 3, a sample from $[\mathbf{x}_t|\mathbf{z}_{1:t}, \boldsymbol{\theta}_{1:t}^{(i)}]$ needs to be computed for each particle with significantly nonzero weight $w_t^{(i)}$. This sample can be obtained using a special case of the Gibbs sampler in Algorithm 2 with $\boldsymbol{\theta}_t = \boldsymbol{\theta}_t^{(i)}$ held fixed (i.e., skipping Step 3(c)). In settings for which the Laplace approximation works well (see Section 3.3), this Gibbs sampler could consist of only a single iteration, in which $\mathbf{y}_t$ is first sampled from the Gaussian approximation to its FCD used in the Laplace approximation, and then $\mathbf{x}_t$ is sampled as in Step 3(a) of Algorithm 2. A PEnKF using this Laplace approach is illustrated for Poisson data in Section 5.2.

Algorithm 3 is, in principle, also applicable when $\boldsymbol{\theta}_t \equiv \boldsymbol{\theta}$ is constant over time (i.e., the prior distribution p_t is a point mass at the previous value). However, in this case the proposal distribution also has to be a point mass, which means that over time there will be fewer and fewer unique particles. To avoid this, the particles can be resampled based on a kernel density estimate of the filtering distribution (Liu and West, 2001; Frei and Künsch, 2012), but then the state ensemble has to be propagated (and updated) again based on the newly sampled parameter values.

As described in Section 3.1, in the case of forecast independence, the likelihood can be computed based on the entire filtering ensemble from $t-1$, and so we can set $N = 1$ and carry out Algorithm 3 with a single ensemble instead of an ensemble of ensembles. This important special case is described in Algorithm 1 of Frei and Künsch (2012). If $\mathcal{M}_t$ does not depend on $\boldsymbol{\theta}_t$ but is computationally expensive, it can then be advantageous to "oversample" the parameters, so that there are more parameter particles than state ensemble members. In Example 4, exploiting forecast independence would allow setting $N = 1$ and using a single $\widetilde{\boldsymbol{\mu}}$ and $\widetilde{\boldsymbol{\Sigma}}$ of based on the entire ensemble $\widetilde{\mathbf{x}}^{(1:M,1)}$.

4.2 Extended ensemble Kalman smoothers

We now discuss extended EnKS methods for HSSMs when smoothing inference is desired; that is, when interest is in $[\mathbf{x}_{1:T}, \boldsymbol{\theta}_{1:T}|\mathbf{z}_{1:T}]$ for some fixed time period $\{1, \ldots, T\}$.

4.2.1 Smoothing inference for high-dimensional parameter vectors

If the parameters (and $\mathbf{y}_t$) are high-dimensional but they can be efficiently sampled from their full-conditional distributions (FCDs), a Gibbs-EnKS can be applied:

Algorithm 4: Gibbs ensemble Kalman smoother (GEnKS)

1. Initialize $\mathbf{y}_{1:T|T}$ and $\boldsymbol{\theta}_{1:T|T}$.
2. Iterate between following steps until convergence:
   (a) Obtain samples $\mathbf{x}_{1:T|T}^{(1:N)}$ from $[\mathbf{x}_{1:T}|\boldsymbol{\theta}_{1:T|T}, \mathbf{y}_{1:T|T}]$ using the EnKS (Algorithm 1), and set $\mathbf{x}_{1:T|T} = \mathbf{x}_{1:T|T}^{(j)}$ for a j sampled uniformly at random from $\{1, \ldots, N\}$.
   (b) Draw a sample $\mathbf{y}_{1:T|T}$ from $[\mathbf{y}_{1:T}|\mathbf{z}_{1:T}, \mathbf{x}_{1:T|T}, \boldsymbol{\theta}_{1:T|T}]$.
   (c) Draw a sample $\boldsymbol{\theta}_{1:T|T}$ from $[\boldsymbol{\theta}_{1:T}|\mathbf{x}_{1:T|T}, \mathbf{y}_{1:T|T}, \mathbf{z}_{1:T}]$.

The GEnKS can be interpreted as an alternative version of the Gibbs-KF of Carter and Kohn (1994) or the particle-Gibbs algorithm of Andrieu et al. (2010), where the Kalman or particle filter, respectively, is replaced by the EnKF. This allows the GEnKS to be applied to high-dimensional, nonlinear SSMs. See Section 5.3 for a numerical comparison.

4.2.2 Smoothing inference for low-dimensional parameters

If $\boldsymbol{\theta}_{1:T}$ cannot be sampled efficiently from its FCD in Step 1(c) of Algorithm 4 but $\boldsymbol{\theta}_t$ is low-dimensional, the following smoothing algorithms could be applied.

Static parameters First, consider the case of a static (temporally constant) low-dimensional parameter vector. In this case, it is possible to apply a Metropolis-Hastings-EnKS (MHEnKS). This algorithm consists of proposing a new parameter value, $\boldsymbol{\theta}^P$ say, and then accepting this value with a probability depending on the likelihood

$$[\mathbf{z}_{1:T}|\boldsymbol{\theta}] = \prod_{t=1}^{T} [\mathbf{z}_t|\mathbf{z}_{1:t-1}, \boldsymbol{\theta}] = \prod_{t=1}^{T} \mathcal{L}_t^{\mathrm{Z}}(\mathbf{z}_t|\boldsymbol{\theta}, \widetilde{\mathbf{x}}_{t|t-1}^{(1:N)}) \tag{18}$$

for $\boldsymbol{\theta} = \boldsymbol{\theta}^P$, where the terms in the product on the right-hand side are the integrated likelihoods at time t in (17), which can be obtained sequentially by carrying out an EnKF and computing the integrated likelihood at each time t. Then, we can obtain samples from $[\mathbf{x}_{1:T}|\mathbf{z}_{1:T}, \boldsymbol{\theta}_{1:T}^{(i)}]$ for each (thinned) posterior sample $\boldsymbol{\theta}_{1:T}^{(i)}$ using a special case of Algorithm 4, in which we hold $\boldsymbol{\theta}_{1:T} = \boldsymbol{\theta}_{1:T}^{(i)}$ fixed and skip Step 2(c). If the evolution is linear, it is also possible to reuse the filtering ensemble $\mathbf{x}_{T|T}^{(1:N)}$ already obtained in the EnKF by applying the EnKS from Stroud et al. (2010).

The MHEnKS can be interpreted as an alternative version of the particle marginal Metropolis-Hastings algorithm of Andrieu et al. (2010), where the particle likelihood in (15) is replaced by the EnKF likelihood in (16). The MHEnKS should hence scale to higher state dimensions.

Time-varying parameters If $\boldsymbol{\theta}_t$ varies over time, inference becomes more challenging, in that even if each $\boldsymbol{\theta}_t$ consists of a small number of elements, the combined vector of unknown parameters at all time points, $\boldsymbol{\theta}_{1:T}$, can still be rather high-dimensional if T is large. Hence, for a low-dimensional time-varying parameter vector, we propose the following particle-EnKS:

Algorithm 5: Particle ensemble Kalman smoother (PEnKS)

1. Obtain smoothing trajectories $\boldsymbol{\theta}_{1:T}^{(1:M)}$ and corresponding weights $w_T^{(1:M)}$ by running the PEnKF in Algorithm 3 up to time T, and saving (and potentially resampling) the trajectories $(\boldsymbol{\theta}_{1:t}^{(i)}, \mathbf{x}_{1:t|t}^{(i,1:N)})$ at each time t (Kitagawa, 1996).

2. If desired, improve degeneracy by obtaining new state trajectories $\mathbf{x}_{1:T|T}^{(i,1:N)}$ from $[\mathbf{x}_{1:T}|\mathbf{y}_{1:T}, \boldsymbol{\theta}_{1:T}^{(i)}]$ for each particle trajectory $\boldsymbol{\theta}_{1:T}^{(i)}$ with significantly nonzero weight $w_T^{(i)}$, using a special case of Algorithm 4, in which we hold $\boldsymbol{\theta}_{1:T} = \boldsymbol{\theta}_{1:T}^{(i)}$ fixed and skip Step 2(c).

Note that this algorithm might perform poorly when T is large, because of degeneracy problems at "early" time points (e.g., Doucet et al., 2000b). For Step 1, we can obtain the same simplifications in the case of forecast independence as for the "stand-alone" PEnKF described in Section 4.1.2. For Step 2, we can again re-use state ensembles in the case of linear evolution using the EnKS procedure of Stroud et al. (2010).

4.2.3 A more general MCMC-EnKS

An extension of the Gibbs-EnKS in Algorithm 4, which we will call an MCMC-EnKS, is possible as long as $\boldsymbol{\theta}_t$ can be split into a set of parameters $\boldsymbol{\theta}_{1:T} = (\boldsymbol{\theta}_{1:T}^{(1)\prime}, \ldots, \boldsymbol{\theta}_{1:T}^{(J)\prime})'$, where each $\boldsymbol{\theta}_{1:T}^{(j)}$ is itself either low-dimensional or has a closed-form full-conditional distribution, and can be updated as in the Metropolis-Hastings-, Gibbs-, or particle-EnKS. This allows the fitting HSSMs with a variety of unknown parameters.

4.3 Properties of the extended EnKF algorithms

4.3.1 Convergence

For linear Gaussian SSMs, the EnKF converges to the Kalman filter as $N \to \infty$ (e.g., Mandel et al., 2011), as long as $\widehat{\boldsymbol{\Sigma}}_{t|t-1}$ is a consistent estimator of the forecast covariance matrix (i.e., the tapering range in Section 2.2 tends to infinity as $N \to \infty$). In that setting, the EnKF likelihood $\mathcal{L}_t^{\mathrm{E}}$ also converges to the true likelihood at each time t. Combining this with convergence results for MCMC (e.g., Robert and Casella, 2004) and SMC (e.g., Doucet et al., 2001) methods under mild regularity conditions, it is clear that the PEnKF, GEnKS, MHEnKS, and PEnKS all converge to the true filtering or smoothing distributions in a HSSM with linear evolution operators $\mathcal{M}_t$, as $N \to \infty$ and M or the number of MCMC iterations increase. This is not exactly the case for the GEnKF, in that the marginal (integrating over $\boldsymbol{\theta}_t$ and $\mathbf{y}_t$) filtering distribution of $\mathbf{x}_t$ and hence the forecast distribution of $\mathbf{x}_{t+1}$ are not exactly Gaussian anymore, but the GEnKF treats them as such.

For SSMs with nonlinear evolution, the EnKF essentially converges to the best filter among those which update the state linearly based on the data at each time point (e.g., Le Gland et al., 2011; Law et al., 2016). As argued in Section 3.2.2, the EnKF likelihood then converges

to the (proper) Dawid-Sebastiani score. Thus, for nonlinear evolution the extended EnKF and EnKS techniques will not produce exact inference, but the resulting samples of the state and parameters can still provide useful approximations to the desired distributions in many settings.

For the high-dimensional applications of interest here, only small ensembles are computationally feasible, and so the algorithms' performance for small N and M is much more important than their asymptotic properties for N and M tending to infinity. Our algorithms will tend to perform well for HSSMs for which the embedded SSM (2)–(3) is well suited for inference using the EnKF and EnKS. This is the case for many real-world applications (e.g., Houtekamer and Mitchell, 2005; Bonavita et al., 2010; Houtekamer et al., 2014). We have also seen in Section 3.2, that the EnKF approximation to the likelihood can be much less variable than the particle approximation, which should lead to better algorithms. Further numerical comparisons are provided in Section 5 below.

4.3.2 Computational cost

The EnKF requires storing, propagating, and operating on N vectors (ensemble members) of length n, and the computational complexity of each EnKF update is $\mathcal{O}(nN^2)$ for most EnKF variants (e.g., Tippett et al., 2003). The computations at each time step are highly parallelizable.

For our extended EnKF algorithms, the EnKF generally has to be carried out several times. However, often only a small number of iterations or particles is necessary to obtain a reasonable approximation to the desired distribution (see Section 5 below). In addition, in the case of forecast independence for the GEnKF and PEnKF, only the update step has to be carried out several times, while the forecast step only has to be carried out for a single ensemble. Thus, considering that the most expensive part of the EnKF is often the forward propagation using an expensive evolution model $\mathcal{M}_t$, the increased computational cost of the GEnKF or PEnKF relative to the EnKF can be negligible.

5 Examples and numerical comparisons

The extended EnKF methods described in Section 4 are very general and can be applied in many situations. We now highlight several interesting examples.

5.1 Examples of non-Gaussian observation models

First, we consider some examples of non-Gaussian observations that lead to closed-form full conditionals and forecast independence and can thus be handled using our GEnKF (Algorithm 2).

5.1.1 Normal scale mixtures

Normal scale mixtures (Andrews and Mallows, 1974; West, 1987) are an important class of distributions that fit in our HSSM framework. Specifically, assume that the observation model is given by $\mathbf{y}_t|\mathbf{x}_t, \boldsymbol{\theta}_t \sim \mathcal{N}_{m_t}(\mathbf{H}_t\mathbf{x}_t, \mathbf{R}_t)$, where now $\mathbf{R}_t = \text{diag}(\sigma^2_{t,1}\theta_{t,1}, \ldots, \sigma^2_{t,m_t}\theta_{t,m_t})$ is diagonal, the $\sigma^2_{t,l}$ are fixed, $\boldsymbol{\theta}_t = (\theta_{t,1}, \ldots, \theta_{t,m_t})'$, $\theta_{t,l} \overset{ind.}{\sim} p_{t,l}(\theta_{t,l})$ does not depend on $\boldsymbol{\theta}_{t-1}$ (i.e., the $\boldsymbol{\theta}_t$ are

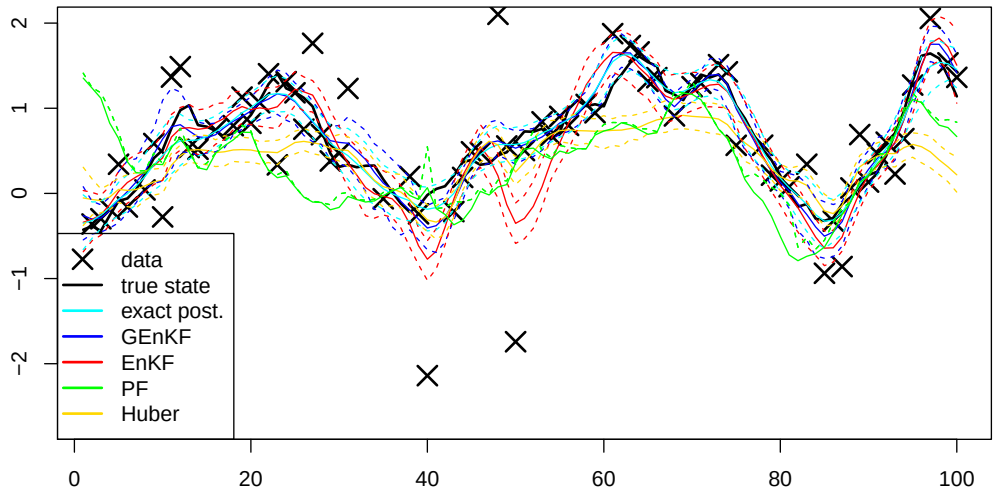

Figure 3: For data with t_2 errors on a one-dimensional spatial domain, posterior means (solid lines) and pointwise posterior .1- and .9-quantiles (dashed lines) of the state as obtained by various methods

independent over time), each $p_{t,l}$ is a distribution with support on the positive real line, and the transformation layer is a point mass such that $\mathbf{z}_t = \mathbf{y}_t$.

This model can be easily fitted using a special case of the GEnKF in Algorithm 2, where we can skip Step 3(b) because $\mathbf{z}_t = \mathbf{y}_t$. In Step 3(c), the parameters can be sampled independently because $\mathbf{R}_t$ is diagonal. In addition, if the $\theta_{t,l}$ follow an inverse-Gamma prior distribution, the full conditional distributions of the parameters are also inverse-Gamma distributions. Specifically, for $\theta_{t,l} \overset{ind}{\sim} \mathcal{IG}(a, b)$, Step 3(c) consists of sampling from

$$\theta_{t,l}|\mathbf{y}_{1:t}, \mathbf{x}_{t|t}^{(i)} \overset{ind}{\sim} \mathcal{IG}\Big(a + \tfrac{1}{2}, b + \Big(\tfrac{y_{t,l} - (\mathbf{H}_t \mathbf{x}_{t|t}^{(i)})_l}{\sigma_{t,l}}\Big)^2/2\Big). \tag{19}$$

A robust ensemble Kalman filter (REnKF) One example of a scale mixture of normals is the t-distribution, which is similar to a normal distribution but has heavier tails and is thus more robust to outliers. Hence, when outliers are present in a given dataset, it can be advantageous to assume the observation model to have noise according to the t-distribution with few degrees of freedom, κ. Specifically, assuming $\theta_{t,l} \overset{ind.}{\sim} \mathcal{IG}(\kappa/2, \kappa/2)$ is equivalent to $(z_{t,l} - y_{t,l})/\sigma_{t,l} \overset{ind.}{\sim} t_\kappa$, meaning that the observation noise is independently t-distributed with κ degrees of freedom. An example of such data with heavy-tailed noise from a t distribution with $\kappa = 2$ degrees of freedom is shown in Figure 3.

Hence, a robust EnKF is given by Algorithm 2, where Step 3(c) consists of sampling from (19) with $a = b = \kappa$ (and skipping Step 3(b) because $\mathbf{z}_t = \mathbf{y}_t$). If κ is unknown with prior $[\kappa]$, it could be sampled in a collapsed Gibbs sampling scheme at the beginning of Step 3(c) (prior to updating $\boldsymbol{\theta}_t$) by sampling from $[\kappa|\mathbf{y}_t, \mathbf{x}_{t|t}^{(i)}] \propto [\kappa] \prod_{l=1}^{m_t} t_\kappa\big(\frac{y_{t,l} - (\mathbf{H}_t \mathbf{x}_{t|t}^{(i)})_l}{\sigma_{t,l}}\big)$. Note that a recently proposed iterative Huber-EnKF (Rao et al., 2015) seems to be similar in spirit to our robust EnKF, but it uses the Huber norm in place of the loss function implied by a t-distribution.

5.1.2 Threshold observation models

Data assimilation for non-continuous observations distributions is particularly challenging, in that the EnKF and its transformation- or quantile-based extensions (e.g., Anderson, 2009; Amezcua and Van Leeuwen, 2014) are often problematic. Two interesting examples of this are binary data and right-skewed data with a point mass at zero (e.g., hourly rainfall amounts). (We also consider count data in Section 5.2.) Binary and rainfall data fit into our HSSM framework in (1)–(4), by assuming that the transformation distribution f_t is determined by transformations of the form $z_{t,l} = g_t(y_{t,l}; \boldsymbol{\theta}_t)$, $l = 1, \ldots, m_t$, where $g_t(\cdot)$ is a deterministic function containing an indicator function, and we assume $\mathbf{R}_t = \text{diag}(\sigma^2_{t,1}, \ldots, \sigma^2_{t,m_t})$ for simplicity.

Inference for such threshold models can be carried out using the GEnKF in Algorithm 2, where Step 3(b) consists of sampling $y^{(i)}_{t,l}$ from $[y_{t,l}|z_{t,l}, \mathbf{x}^{(i)}_{t|t}, \boldsymbol{\theta}^{(i)}_t]$, independently for $l = 1, \ldots, m_t$ (due to $\mathbf{R}_t$ being diagonal). Step 3(b) can be skipped if there are no unknown parameters. We now give specifics for the examples of binary and rainfall-type data.

Binary observations Binary observations follow a Bernoulli distribution, $z_{t,l}|\mathbf{x}_t, \boldsymbol{\theta}_t \sim B\big(\Phi((\mathbf{H}_t\mathbf{x}_t)_l/\sigma_{t,l})\big)$, also called a probit model for $\sigma_{t,l} = 1$, if we set

$$z_{t,l} = g(y_{t,l}; \boldsymbol{\theta}_t) = \begin{cases} 1, & y_{t,l} > 0, \\ 0, & y_{t,l} \leq 0. \end{cases}$$

For Step 3(b) in Algorithm 2, we then have $[y_{t,l}|z_{t,l}, \mathbf{x}_t, \boldsymbol{\theta}] = \mathcal{N}^{+/-}\big((\mathbf{H}_t\mathbf{x}^{(i)}_t)_l, \sigma^2_{t,l}\big)$, where $\mathcal{N}^{+/-}$ is a truncated normal distribution restricted to the positive or negative real line for $z_{t,l}$ equal to one or zero, respectively. In this example we have no unknown parameters, $\boldsymbol{\theta}_t$.

Rainfall observations For the rainfall example, assume that $z_{t,l}$ is the lth rainfall measurement observed at time t. Then $z_{t,l}$ should be modeled as a mixture of a positive, right-skewed distribution and a point mass at zero. If $y_{t,l}$ is a latent Gaussian random variable, such a distribution can be obtained by setting

$$z_{t,l} = g(y_{t,l}; \boldsymbol{\theta}_t) = \begin{cases} y^{\kappa_t}_{t,l}, & y_{t,l} > 0 \\ 0, & y_{t,l} \leq 0 \end{cases}$$

for some $\kappa_t > 1$ that determines the skewness (e.g., Sigrist et al., 2012). An example of data generated from this observation model is shown in Figure 4.

For Step 3(b) in Algorithm 2, we need to sample $[y_{t,l}|z_{t,l}, \mathbf{x}_t, \kappa_t]$, which is now a point mass at $z^{1/\kappa_t}_{t,l}$ if $z_{t,l} > 0$, and $\mathcal{N}^{-}\big((\mathbf{H}_t\mathbf{x}^{(i)}_t)_l, \sigma^2_{t,l}\big)$ otherwise (i.e., a truncated normal distribution on the negative real line). If κ_t is unknown and independent over time with prior $\kappa_t \sim p_\kappa(\kappa_t)$, we can (before updating $\mathbf{y}_t$) sample κ_t from $[\kappa_t|\mathbf{z}_t, \mathbf{x}_t] \propto p_\kappa(\kappa_t) \prod_{l:z_{t,l}>0} \mathcal{N}(z^{1/\kappa_t}_{t,l}|(\mathbf{H}_t\mathbf{x}_t)_l, \sigma^2_{t,l}) J_l(\kappa_t)$, where $J_l(\kappa_t) = z^{(1/\kappa_t)-1}_{t,l}/\kappa_t$ is the Jacobian for the transformation from $z_{t,l}$ to $z^{1/\kappa_t}_{t,l}$, using a Metropolis-Hastings update.

5.1.3 Simulation study for non-Gaussian observations

We compared our new GEnKF to several alternatives in a simulation study for three non-Gaussian observation scenarios considered here: heavy-tailed observation noise from a t distribution with $\kappa = 2$ degrees of freedom (see Figure 3), and rainfall data with the true $\kappa = 3$

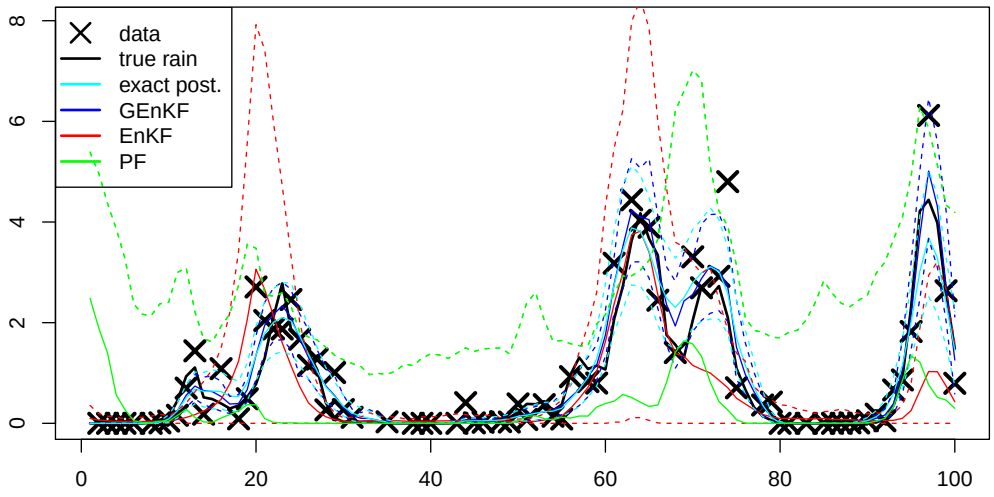

Figure 4: For rainfall amounts with known $\kappa = 3$ on a one-dimensional spatial domain, posterior means (solid lines) and pointwise posterior .1- and .9-quantiles (dashed lines) of the state as obtained by various methods

assumed known (see Figure 4) or unknown. Because non-Gaussian observations are not directly related to the evolution model and the forecast step, we focus on the update step at a single time point t.

We simulated 100 true state vectors of size $n = 100$ on a one-dimensional spatial domain $[1, 100]$ with $m_t = 75$ randomly chosen observation locations. The state is simulated from a normal distribution with known true mean 0.2 (constant over space) and known true covariance matrix based on a powered exponential covariance function with power 1.8 and scale parameter 10, same for all three simulation scenarios. Also, we assumed $\sigma_{t,l} \equiv 0.2$ to be known for all examples.

The GEnKF was compared to a standard particle filter (importance sampler), and to the true posterior distribution obtained using an MCMC that was run for 1,000 iterations after 500 burn-in iterations (which was only possible here because of the low-dimensional examples and the focus on the update at a single time point). We also compared our GEnKF to a commonly used extension of the EnKF update (Anderson, 2003). For an observation equation of the form $\mathbf{y}_t = \mathcal{H}_t(\mathbf{x}_t, \mathbf{v}_t)$, the pseudo-observation in (12) is then replaced by $\widetilde{\mathbf{y}}_t^{(j)} = \mathcal{H}_t(\mathbf{x}_{t|t-1}^{(j)}, \mathbf{v}_t^{(j)})$, and the Kalman gain is estimated as $\widehat{\mathbf{K}}_t = \widehat{\mathbf{C}}_{xy}\widehat{\mathbf{C}}_{yy}^{-1}$, where $\widehat{\mathbf{C}}_{yy}$ and $\widehat{\mathbf{C}}_{xy}$ are (potentially regularized versions of) the empirical covariance matrix of $\widetilde{\mathbf{y}}_t^{(1:N)}$ and the cross-covariance matrix between $\mathbf{x}_{t|t-1}^{(1:N)}$ and $\widetilde{\mathbf{y}}_t^{(1:N)}$, respectively. For the rainfall data with unknown κ, we used state augmentation. For the heavy-tailed data, we also considered the robust Huber-EnKF proposed in Roh et al. (2013). The Huber-EnKF requires a cutoff or clipping value for each observation location, which we selected by sampling 1000 times from the true forecast distribution based on an update consisting of a single observation at the respective location.

For the standard EnKF, the Huber-EnKF, and the particle filter (PF), we used the same prior ensemble of size $M = 100$, simulated from the true prior distribution. For our GEnKF we only used a subset of this ensemble of size $M = 30$, and only one (for rainfall with known κ) or three (for the other two scenarios) Gibbs iterations. For the EnKF and GEnKF, we used the same Wendland taper function with tapering length 20.

The results are shown in Figures 3 and 4 and in Table 2. We computed both the mean

	Heavy-tailed		Rainfall (κ known)		Rainfall (κ unknown)	
	MSPE	CRPS	MSPE	CRPS	MSPE	CRPS
exact	0.175	0.096	0.452	0.146	0.450	0.146
GEnKF	**0.195**	**0.107**	**0.470**	**0.154**	**0.895**	**0.232**
EnKF	0.289	0.154	13.907	3.848	>100	>100
PF	0.764	0.582	2.033	1.037	4.289	1.861
Huber	0.629	0.405				

Table 2: Simulation results for non-Gaussian observations. Heavy-tailed data are described in Section 5.1.1, and binary and rainfall data are described in Section 5.1.2. MSPE = mean squared prediction error; CRPS = continuous ranked probability score; exact = posterior distribution obtained using MCMC; PF = particle filter (importance sampling); Huber = robust Huber-EnKF from Roh et al. (2013)

squared prediction error (MSPE), which evaluates the accuracy of the best point prediction (under squared error loss), and the continous ranked probability score (CRPS), which evaluates the quality of the entire predictive distribution, including its shape and spread (see, e.g., Gneiting and Katzfuss, 2014). For the heavy-tailed observations, we compared the state filtering distribution for the different methods to the true state, whereas for rainfall data we considered predictions of the true rainfall, $x_{t,l}^{\kappa} I(x_{t,l} > 0)$.

The GEnKF outperformed all competitors in all three scenarios, and came relatively close to the true posterior distribution. This is especially remarkable considering the GEnKF consisted of at most three EnKF updates with ensemble size 30, while the competitors each used an ensemble of size 100. Hence, in our simulations, the GEnKF produced the most accurate results at the lowest computational expense.

5.2 Cloud data

We also conducted filtering inference for a cloud-motion dataset (shown in Figure 5a) from Wikle (2002), representing cloud intensities (i.e., counts) at $n = 60$ locations along a spatial transect at $T = 80$ time points. The data roughly follow an overdispersed Poisson distribution, and so in our HSSM we set $f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t) = \prod_{l=1}^{m_t} Poi(z_{t,l}|e^{y_{t,l}})$. We randomly set aside 10% of the observations as test data, so that $m_t \equiv 54$ and $\mathbf{H}_t$ was a subset of the identity. Further, we assumed $\mathbf{R}_t = \sigma_t^2 \mathbf{I}_{m_t}$, where σ_t^2 accounts for overdispersion. The evolution was taken to be linear, $\mathcal{M}_t(\mathbf{x}_{t-1}) = \mathbf{M}_t \mathbf{x}_{t-1}$, where $\mathbf{M}_t$ was a tridiagonal matrix with $\gamma_{1,t}$ on the main diagonal, and $\gamma_{2,t}$ and $\gamma_{3,t}$ on the first diagonals above and below this, respectively. The evolution error was assumed to exhibit spatial dependence in the form of a Matérn covariance with smoothness 1.5, such that the (i,j)th element of $\mathbf{Q}_t$ is given by $(\mathbf{Q}_t)_{i,j} = \tau_t^2 \,\mathrm{Mat}(|i-j|/\lambda_t)$, where $\mathrm{Mat}(d) = (1+\sqrt{3}d)\exp(-\sqrt{3}d)$. The unknown parameters were $\boldsymbol{\theta}_t := (\gamma_{1,t}, \gamma_{2,t}, \gamma_{3,t}, \log\sigma_t, \log\tau_t, \log\lambda_t)'$ with $p_t(\boldsymbol{\theta}_t|\boldsymbol{\theta}_{t-1}) = \mathcal{N}_6(\boldsymbol{\theta}_{t-1}, 0.05^2\mathbf{I}_6)$. For $t = 0$, we set $\boldsymbol{\theta}_0 = (0.3, 0.3, 0.3, \log(0.1), \log(1.5), \log(8))'$, $\boldsymbol{\mu}_{0|0} = (-2)\mathbf{1}_n$, and $(\boldsymbol{\Sigma}_{0|0})_{i,j} = 0.2\,\mathrm{Mat}(|i-j|/5)$. Note that our model addresses the potential extension to time-varying parameters suggested in Wikle (2002).

We considered two methods for filtering inference on the state and parameters: our PEnKF from Section 4.1.2, and an exact approximation of a Rao-Blackwellized particle filter (EARBPF; Johansen et al., 2012). Forecast independence does not hold for all parameters in this example, and so we carried out the full PEnKF from Algorithm 3 with $M = 50$ ensembles of size $N = 30$. For tapering of the forecast covariance matrix (see Section 2.2), we used a Wendland taper with

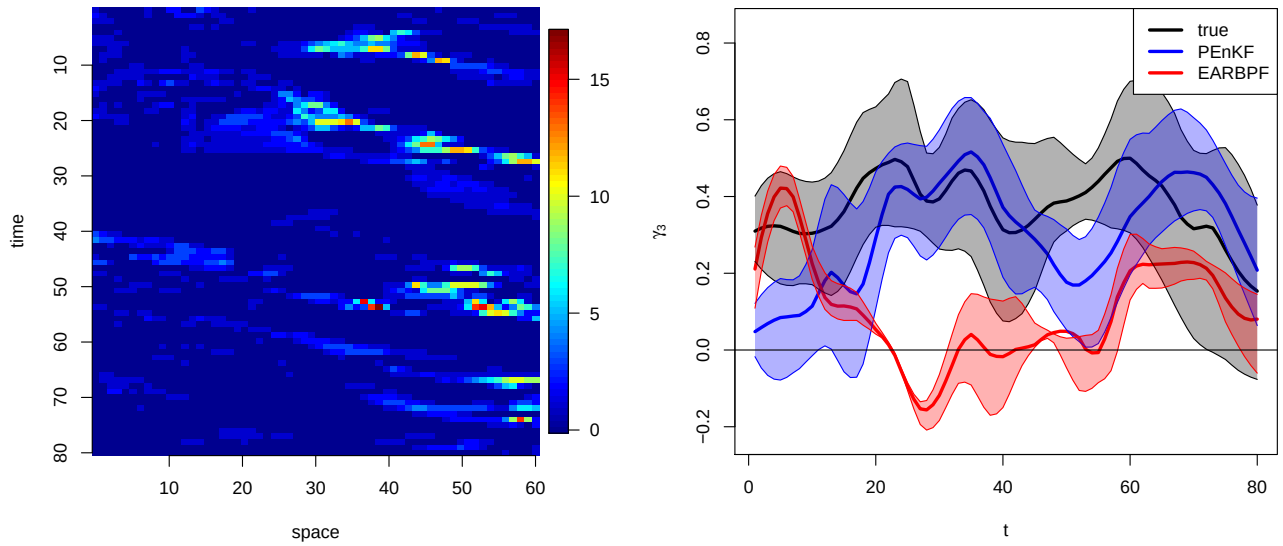

(a) Cloud intensity data

(b) Filtering distributions for $\gamma_{3,t}$

Figure 5: The cloud data from Section 5.2 (Panel (a)), and filtering means and 95% credible intervals (slightly smoothed for clearness) over time for $\gamma_{3,t}$ (Panel (b))

range 8. The integrated likelihood and posterior samples of $\mathbf{x}_t$ in Steps 1(c) and 2, respectively, of Algorithm 3 were obtained using a Laplace approximation (see Section 3.3 and the text below Algorithm 3). For the EARBPF, which is similar to the PEnKF except that it uses local particle filters instead of the EnKF (and the Laplace approximation) to integrate out the state for each parameter particle, we considered one setting with a similar computation time as the PEnKF ($M = 50$ and $N = 30$), and one setting with a much larger computation time ($M = 500$ and $N = 10000$) that should provide a close approximation to and will henceforth be referred to as the true filtering distribution.

	Predictions		EMD to true filtering distributions					
	MSPE	CRPS	γ_1	γ_2	γ_3	$\log\sigma$	$\log\tau$	$\log\lambda$
truth	0.73	0.25						
PEnKF	**0.75**	**0.25**	**0.25**	**0.16**	**0.19**	**0.70**	**0.23**	**0.32**
EARBPF	1.26	0.33	0.27	0.25	0.24	0.71	0.32	0.37

Table 3: Mean squared prediction error (MSPE) and continuous ranked probability score (CRPS) for prediction of test data, and earth-mover's distance (EMD) to the true parameter posteriors, for the cloud data in Section 5.2

The results, averaged over ten different initial ensembles $\mathbf{x}_{0|0}^{(1:M,1:N)}$, are shown in Table 3. In terms of filtering predictions of the test data, the PEnKF performed almost as well as the true filtering distributions, and considerably better than the EARBPF. The PEnKF approximation to the filtering distributions of the six parameters in $\boldsymbol{\theta}_t$ is also closer than the EARBPF to the true distributions in terms of earth-mover's distance (computed using the R package by Urbanek and Rubner, 2012) averaged over time. As an example, Figure 5b shows the filtering distributions (for one initial ensemble) over time of the parameter $\gamma_{3,t}$, which can be viewed as quantifying the amount of movement of the cloud intensities "one pixel to the right" from one time point to the next. This rightward drift, which is clearly present in the data (see Figure 5a),

is picked up fairly quickly by the true and PEnKF filtering distributions, while the EARBPF approximation even implies a negative $\gamma_{3,t}$ between $t = 25$ and $t = 30$.

5.3 Smoothing inference for the Lorenz-96 model

Finally, we considered smoothing inference for simulated data based on the popular nonlinear Lorenz-96 model (Lorenz, 1996). Lorenz-96 mimics advection at $n = 40$ equally-spaced locations along a latitude circle on the globe. The evolution over time is determined by the differential equations

$$\frac{dx_{t,i}}{dt} = (x_{t,i+1} - x_{t,i-2})x_{t,i-1} - x_{t,i} + F, \qquad i = 1, \ldots, 40, \tag{20}$$

with cyclic boundary conditions (i.e., $x_{t,41} = x_{t,1}$, and so forth). Let $\text{Lor}_{8,0.2}$ denote the evolution operator implied by (20) with forcing parameter $F = 8$ and time step $\delta = 0.2$, which leads to strong nonlinearities and non-Gaussian forecast distributions. The Lorenz-96 equations were solved numerically using an Euler scheme. Let $\boldsymbol{\Sigma}_L$ denote the covariance matrix obtained by a long run of the $\text{Lor}_{8,0.2}$ model.

We simulated data from our HSSM (1)–(4) for $t = 1, \ldots, T$ with $T = 10$, $f_t(\mathbf{z}_t|\mathbf{y}_t, \boldsymbol{\theta}_t) = \delta_{\mathbf{y}_t}(\mathbf{z}_t)$ (i.e., $\mathbf{z}_t = \mathbf{y}_t$), $\mathcal{M}_t(\mathbf{x}_{t-1}) = \theta\,\text{Lor}_{8,0.2}(\mathbf{x}_{t-1})$, $\mathbf{Q}_t = 0.2\boldsymbol{\Sigma}_L$, $\mathbf{H}_t = \mathbf{I}_n$, $\mathbf{R}_t = \mathbf{I}_n$, $\boldsymbol{\mu}_{1|0} = \mathbf{0}$, $\boldsymbol{\Sigma}_{1|0} = \boldsymbol{\Sigma}_L$. We assumed θ to be the unknown (static) parameter with prior $\mathcal{N}(\mu_\theta, \sigma_\theta^2)$, where $\mu_\theta = 0.8$ and $\sigma_\theta = 0.2$.

We simulated 100 such datasets, with the true θ drawn for each simulation from the prior. For each dataset, we ran three different methods to obtain approximations to the smoothing distribution $[\theta, \mathbf{x}_{1:T}|\mathbf{y}_{1:T}]$ using a starting value of $\theta = 0.5$. The first method considered was the EnKS in Algorithm 1 with N =1,000 and state augmentation; that is, θ was treated as part of the state vector $\mathbf{x}_t$ with artificial evolution distribution $\theta_t|\theta_{t-1} \sim \mathcal{N}(\theta_{t-1}, 0.1^2)$. We also considered the GEnKS in Algorithm 4. For both EnKS algorithms with used a taper over time and space given by a Wendland correlation function with radius 3 and 8, respectively. Finally, we considered the particle Gibbs sampler (Andrieu et al., 2010, Sect. 2.4.3). For both the GEnKS and the particle Gibbs sampler, we used an ensemble size of $N = 50$ and 100 Gibbs iterations, the first 20 of which were considered burn-in. For the GEnKS in Algorithm 4, we can skip sampling $\mathbf{y}_{1:T}$ (Step 2(b)) because here $\mathbf{y}_t = \mathbf{z}_t$ was given by the data. The full-conditional distribution (FCD) of the static parameter θ in Step 2(c) is available in closed form:

$$\theta|\mathbf{x}_{1:T}, \mathbf{y}_{1:T} \sim \mathcal{N}\Big(\tfrac{\mu_\theta/\sigma_\theta^2 + \sum_{t=2}^T \widetilde{\mathbf{x}}_t'\mathbf{Q}_t^{-1}\mathbf{x}_t}{1/\sigma_\theta^2 + \sum_{t=1}^T \widetilde{\mathbf{x}}_t'\mathbf{Q}_t^{-1}\widetilde{\mathbf{x}}_t}, \tfrac{1}{1/\sigma_\theta^2 + \sum_{t=2}^T \widetilde{\mathbf{x}}_t'\mathbf{Q}_t^{-1}\widetilde{\mathbf{x}}_t}\Big),$$

where $\widetilde{\mathbf{x}}_t = \text{Lor}_{8,0.2}(\mathbf{x}_{t-1})$. The same distribution is used to update θ in the particle Gibbs sampler, in which a particle filter is used to obtain the sample from the FCD of the state $\mathbf{x}_{1:T}$ given the data and θ.

Due to the nonlinear evolution, it is not possible to obtain the true smoothing distribution of the states and parameter in any reasonable computation time. Hence, we compared the three methods using the same proper scoring rules as in Section 5.1.3, averaged over the 100 simulated datasets. The results are summarized in Table 4. Our GEnKS algorithm performed best both in terms of inference on the parameter and on the states. The state augmentation method failed, leading to very high (i.e., bad) scores for the posterior on θ. The particle Gibbs sampler produced worse inference on θ than simply using the prior distribution (i.e., completely ignoring any information in the data). Some of the results are also summarized in Figure 6, which shows

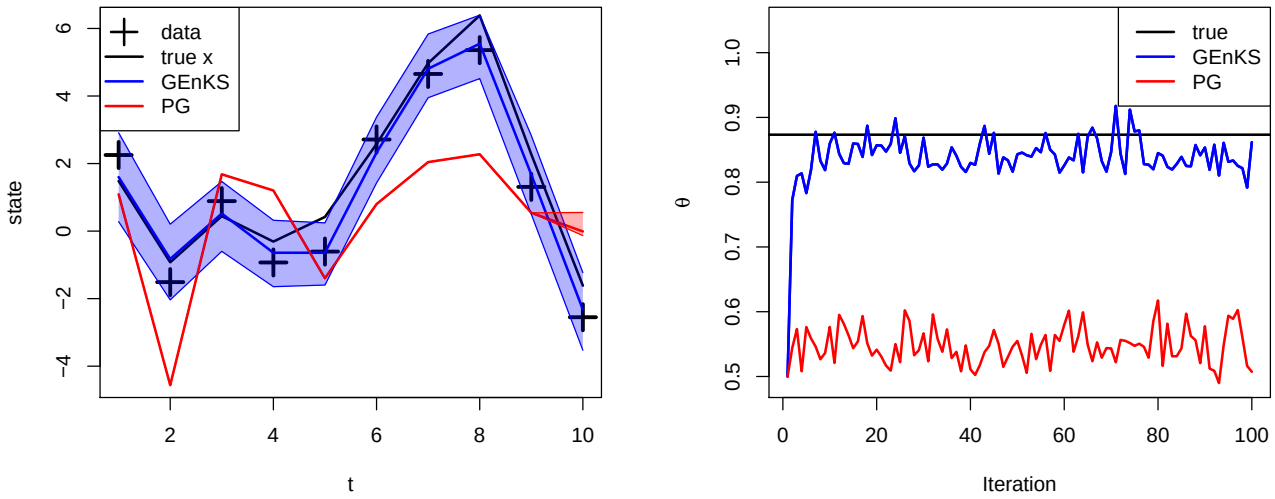

(a) State at location 1 over time (i.e., $\mathbf{x}_{1:T,1}$)

(b) Trace plots for θ

Figure 6: For one of the Lorenz-96 simulations in Section 5.3, true values and posterior distributions (posterior means and pointwise 80% prediction intervals) of the state at a single location (Panel (a)), and trace plots for θ (Panel (b))

the collapse of the particle Gibbs procedure for inference on $\mathbf{x}_{1:T}$ and its non-convergence in terms of inference on θ.

	Parameter θ		State $\mathbf{x}_{1:T}$	
	MSPE	CRPS	MSPE	CRPS
GEnKS	**0.002**	**0.024**	**0.710**	**0.478**
EnKS+SA	>100	>100	0.914	0.540
Particle Gibbs	0.111	0.262	12.380	2.495
Prior	0.042	0.118		

Table 4: Mean squared prediction error (MSPE) and continuous ranked probability score (CRPS) for smoothing inference on the parameter θ and states $\mathbf{x}_{1:T}$ for the Lorenz-96 simulations in Section 5.3

Note that while we considered θ to be the only unknown parameter here for clearness of exposition, our GEnKS should also work well for unknown parameters in $\mathbf{H}_t$ and scalar multiplicative parameters in $\mathbf{Q}_t$ and $\mathbf{R}_t$, in that closed-form FCDs are available in both of these cases if a normal and inverse-gamma priors, respectively, are assumed.

6 Discussion

We have presented a class of extended ensemble Kalman filters and smoothers for inference in high-dimensional and potentially nonlinear state-space models. These algorithms consist of existing Bayesian inference techniques, but carry out inference on or integrate out the high-dimensional state vector using the ensemble Kalman filter (EnKF). Extended EnKFs inherit their properties from the EnKF, in that they are approximate, scale to high dimensions, and generally work well if an EnKF is well suited and tuned for inference on the state vector with all parameters known in the application at hand. While we have only considered a stochastic

EnKF with simultaneous updates for the entire state vector, we can easily replace this with deterministic or serial versions of the EnKF (e.g., Tippett et al., 2003; Hunt et al., 2007).

We have shown that the extended EnKFs can outperform existing approaches even for the moderate-dimensional examples considered here. Due to the scalability and success of the EnKF in very high dimensions, our extended EnKFs should also work well in truly high-dimensional problems. Efficient and parallel existing implementations such as the data assimilation research testbed, or DART (Anderson et al., 2009), should be suitable for relatively straightforward extensions as proposed here.

Acknowledgments

Katzfuss' research was partially supported by National Science Foundation (NSF) Grant DMS–1521676 and NSF CAREER Grant DMS–1654083. Wikle acknowledges the support of NSF Grant SES–1132031, funded through the National Science Foundation Census Research Network program. We would like to thank Fuqing Zhang and Amir Nikooienejad for helpful comments and suggestions.

A Proofs

Proof of Proposition 1. For notational simplicity, we drop the $\widetilde{\ }$ on $\widetilde{\mathbf{x}}^{(j)}$. First, note that $\sum_{i=1}^n (x_i^{(j)} - y_i)^2$ follows a noncentral χ^2-distribution, and so we show that

$$\log \mathcal{L}^{\mathrm{P}}(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x}^{(j)}) = -\tfrac{n}{2}\log(2\pi\theta) - \tfrac{1}{2\theta}\sum_{i=1}^n (x_i^{(j)} - y_i)^2 \xrightarrow{d} \mathcal{N}(na, nb),$$

where $a = -\frac{1}{2}(\log(2\pi\theta) + \kappa + \frac{1}{n}\sum_{i=1}^n y_i^2)$ and $b = 2\kappa(\kappa + \frac{2}{n}\sum_{i=1}^n y_i^2)$. Using properties of the lognormal distribution, we have

$$\mathcal{L}^{\mathrm{P}}(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x}^{(1:N)}) = \tfrac{1}{N}\sum_{j=1}^N \exp(\log \mathcal{L}^{\mathrm{P}}(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x}^{(j)})) \xrightarrow{d} \mathcal{N}(e^{n(a+b/2)}, \tfrac{1}{N}e^{n2(a+b)}),$$

and finally, using the delta-method, we have the asymptotic variance

$$\operatorname{var} \log \mathcal{L}^{\mathrm{P}}(\mathbf{y}|\boldsymbol{\theta}, \mathbf{x}^{(1:N)}) \overset{a}{=} \tfrac{1}{N}e^{n2(a+b)}e^{-2n(a+b/2)} = \tfrac{1}{N}e^{nb} = \tfrac{1}{N}e^{n2\kappa(\kappa+(2/n)\sum_{i=1}^n y_i^2)} = \mathcal{O}(e^n/N).$$

□

Proof of Proposition 2. Due to the diagonal tapering, the covariance matrix in the likelihood is a diagonal matrix with ith element

$$\hat{\sigma}_i^2 = \tfrac{1}{N-1}\sum_{j=1}^N (x_i^{(j)} - \hat{\mu}_i)^2 + \theta \xrightarrow{d} \mathcal{N}(\sigma^2, 2\kappa^2/N),$$

where $\sigma^2 = \kappa + \theta$ and $\hat{\mu}_i = \frac{1}{N}\sum_{j=1}^N x_i^{(j)}$. Because $\operatorname{var}(\hat{\mu}_i^2) = \mathcal{O}(1/N^2)$, the variance of the mean estimator is negligible relative to that of $\hat{\sigma}_i^2$, and we regard $c_i = (y_i - \hat{\mu}_i)^2$ as a constant. Thus, by the delta method,

$$\operatorname{var}(\log \mathcal{L}^{\mathrm{E}}(y_i|\theta, x_i^{(1:N)})) = \operatorname{var}(\tfrac{1}{2}\log(\hat{\sigma}_i^2) + \tfrac{1}{2}\tfrac{c_i}{\hat{\sigma}_i^2}) \overset{a}{=} \tfrac{1}{4}\tfrac{2\kappa^2}{N}(\tfrac{1}{\sigma^2} - \tfrac{c_i}{\sigma^4})^2.$$

Finally, we have $\log \mathcal{L}^{\mathrm{E}}(\mathbf{y}|\theta, \mathbf{x}^{(1:N)}) = \sum_{i=1}^n \log \mathcal{L}^{\mathrm{E}}(y_i|\theta, x_i^{(1:N)})$, and so $\operatorname{var}(\log \mathcal{L}^{\mathrm{E}}(\mathbf{y}|\theta, \mathbf{x}^{(1:N)})) = \mathcal{O}(n/N)$. □

References

Amezcua, J. and Van Leeuwen, P. J. (2014). Gaussian anamorphosis in the analysis step of the EnKF: a joint state-variable/observation approach. *Tellus A*, 66:1–18.

Anderson, J., Hoar, T., Raeder, K., Liu, H., Collins, N., Torn, R., and Avellano, A. (2009). The data assimilation research testbed: A community facility. *Bulletin of the American Meteorological Society*, 90(9):1283–1296.

Anderson, J. L. (2001). An ensemble adjustment Kalman filter for data assimilation. *Monthly Weather Review*, 129:2884–2903.

Anderson, J. L. (2003). A local least squares framework for ensemble filtering. *Monthly Weather Review*, 131:634–642.

Anderson, J. L. (2007a). An adaptive covariance inflation error correction algorithm for ensemble filters. *Tellus A*, 59:210–224.

Anderson, J. L. (2007b). Exploring the need for localization in ensemble data assimilation using an hierarchical ensemble filter. *Physica D*, 230:99–111.

Anderson, J. L. (2009). Ensemble Kalman filters for large geophysical applications. *IEEE Control Systems Magazine*, 29(June):66–82.

Andrews, D. and Mallows, C. (1974). Scale mixtures of normality. *Journal of the Royal Statistical Society, Series B*, 36:99–102.

Andrieu, C., Doucet, A., and Holenstein, R. (2010). Particle Markov chain Monte Carlo. *Journal of the Royal Statistical Society, Series B*, 72:1–33.

Andrieu, C. and Roberts, G. O. (2009). The pseudo-marginal approach for efficient Monte Carlo computations. *Annals of Statistics*, 37(2):697–725.

Bengio, Y. and Delalleau, O. (2009). Justifying and generalizing contrastive divergence. *Neural Computation*, 21(6):1601–1621.

Berzuini, C., Best, N. G., Gilks, W. R., and Larizza, C. (1997). Dynamic conditional independence models and markov chain monte carlo methods. *Journal of the American Statistical Association*, 92(440):1403–1412.

Birgé, L. and Massart, P. (1993). Rates of convergence for minimum contrast estimators. *Probability Theory and Related Fields*, 97:113–150.

Bocquet, M. and Sakov, P. (2014). An iterative ensemble Kalman smoother. *Quarterly Journal of the Royal Meteorological Society*, 140:1521–1535.

Bonavita, M., Torrisi, L., and Marcucci, F. (2010). Ensemble data assimilation with the CNMCA regional forecasting system. *Quarterly Journal of the Royal Meteorological Society*, 136(646):132–145.

Carter, C. and Kohn, R. (1994). On Gibbs sampling for state space models. *Biometrika*, 81(3):541–553.

Chen, R. and Liu, J. S. (2000). Mixture Kalman filters. *Journal of the Royal Statistical Society, Series B*, 62(3):493–508.

Chib, S. and Jeliazkov, I. (2001). Marginal likelihood from the Metropolis-Hastings output. *Journal of the American Statistical Association*, 96(453):270–281.

Cressie, N. and Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. Wiley, Hoboken, NJ.

Dawid, A. P. and Sebastiani, P. (1999). Coherent dispersion criteria for optimal experimental design. *Annals of Statistics*, pages 65–81.

Del Moral, P. (2004). *Feynman-Kac Formulae: Genealogical and Interacting Particle Systems with Applications*. Springer.

DelSole, T. and Yang, X. (2010). State and parameter estimation in stochastic dynamical models. *Physica D*, 239(18):1781–1788.

Douc, R., Cappé, O., and Moulines, E. (2005). Comparison of resampling schemes for particle filtering. *ISPA 2005. Proceedings of the 4th International Symposium on Image and Signal Processing and Analysis, 2005.*, pages 64–69.

Doucet, A., De Freitas, N., and Gordon, N. (2001). *Sequential Monte Carlo Methods in Practice*. Springer, New York, NY.

Doucet, A., de Freitas, N., Murphy, K., and Russell, S. (2000a). Rao-Blackwellised particle filtering for dynamic Bayesian networks. *Proceedings of the Sixteenth Conference on Uncertainty in Artificial Intelligence*, pages 176–183.

Doucet, A., Godsill, S., and Andrieu, C. (2000b). On sequential Monte Carlo sampling methods for Bayesian

filtering. *Statistics and Computing*, 10(3):197–208.

Doucet, A., Pitt, M. K., Deligiannidis, G., and Kohn, R. (2015). Efficient implementation of Markov chain Monte Carlo when using an unbiased likelihood estimator. *Biometrika*, 102(2):295–313.

Evensen, G. (1994). Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *Journal of Geophysical Research*, 99:10143–10162.

Evensen, G. and van Leeuwen, P. J. (2000). An ensemble Kalman smoother for nonlinear dynamics. *Monthly Weather Review*, 128:1852–1867.

Frei, M. and Künsch, H. R. (2012). Sequential state and observation noise covariance estimation using combined ensemble Kalman and particle filters. *Monthly Weather Review*, 140:1476–1495.

Furrer, R. and Bengtsson, T. (2007). Estimation of high-dimensional prior and posterior covariance matrices in Kalman filter variants. *Journal of Multivariate Analysis*, 98(2):227–255.

Furrer, R., Genton, M. G., and Nychka, D. (2006). Covariance tapering for interpolation of large spatial datasets. *Journal of Computational and Graphical Statistics*, 15(3):502–523.

Gelfand, A. E. and Smith, A. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association*, 85(410):398–409.

Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6(6):721–741.

Gneiting, T. and Katzfuss, M. (2014). Probabilistic forecasting. *Annual Review of Statistics and Its Application*, 1(1):125–151.

Goldstein, M. and Wooff, D. (2007). *Bayes Linear Statistics, Theory & Methods*. John Wiley & Sons.

Gordon, N., Salmond, D., and Smith, A. (1993). Novel approach to nonlinear/non-Gaussian Bayesian state estimation. *Radar and Signal Processing, IEE Proceedings F*, 140(2):107–113.

Hartigan, J. (1969). Linear Bayesian methods. *Journal of the Royal Statistical Society, Series B*, pages 446–454.

Hastie, T., Tibshirani, R., and Friedman, J. (2009). *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*. Springer, 2nd edition.

Hinton, G. E. (2002). Training products of experts by minimizing contrastive divergence. *Neural Computation*, 14(8):1771–1800.

Houtekamer, P. L., Deng, X., Mitchell, H. L., Baek, S.-J., and Gagnon, N. (2014). Higher resolution in an operational ensemble Kalman filter. *Monthly Weather Review*, 142(3):1143–1162.

Houtekamer, P. L. and Mitchell, H. L. (2005). Ensemble Kalman filtering. *Quarterly Journal of the Royal Meteorological Society*, 131(613):3269–3289.

Houtekamer, P. L. and Zhang, F. (2016). Review of the ensemble Kalman filter for atmospheric data assimilation. *Monthly Weather Review*, 144(12):4489–4532.

Hunt, B. R., Kostelich, E. J., and Szunyogh, I. (2007). Efficient data assimilation for spatiotemporal chaos: A local ensemble transform Kalman filter. *Physica D: Nonlinear Phenomena*, 230(1-2):112–126.

Jiang, W. and Tanner, M. A. (2008). Gibbs posterior for variable selection in high-dimensional classification and data mining. *Annals of Statistics*, 36(5):2207–2231.

Johansen, A. M., Whiteley, N., and Doucet, A. (2012). Exact approximation of Rao-Blackwellised particle filters. In *IFAC Proceedings Volumes*, volume 45, pages 488–493.

Kalman, R. (1960). A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45.

Katzfuss, M. and Bhattacharya, A. (2015). Comment on article by Dawid and Musio. *Bayesian Analysis*, 10(2):501–504.

Katzfuss, M., Stroud, J. R., and Wikle, C. K. (2016). Understanding the ensemble Kalman filter. *The American Statistician*, 70(4):350–357.

Khare, S., Anderson, J., Hoar, T., and Nychka, D. (2008). An investigation into the application of an ensemble Kalman smoother to high-dimensional geophysical systems. *Tellus*, 60A:97–112.

Kitagawa, G. (1996). Monte Carlo filter and smoother for non-Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics*, 5:1–25.

Kitagawa, G. (1998). A self-organising state-space model. *Journal of the American Statistical Association*, 93(443):1203–1215.

Law, K. J. H., Tembine, H., and Tempone, R. (2016). Deterministic mean-field ensemble Kalman filtering. *SIAM Journal on Scientific Computing*, 38(3):A1251–A1279.

Le Gland, F., Monbet, V., and Tran, V.-D. (2011). Large Sample Asymptotics for the Ensemble Kalman Filter. In Crisan, D. and Rozovskii, B., editors, *The Oxford Handbook of Nonlinear Filtering*, pages 598–634. Oxford University Press.

Lei, J., Bickel, P. J., and Snyder, C. (2010). Comparison of ensemble Kalman filters under non-Gaussianity. *Monthly Weather Review*, 138(4):1293–1306.

Li, C., Jiang, W., and Tanner, M. A. (2013). General oracle inequalities for Gibbs posterior with application to ranking. *Jounal of Machine Learning Research: Workshop and Conference Proceedings 30*, pages 512–521.

Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society, Series B*, 73(4):423–498.

Liu, J. and West, M. (2001). Combined parameter and state estimation in simulation-based filtering. In Doucet, A., de Freitas, J., and Gordon, N., editors, *Sequential Monte Carlo Methods in Practice*. Springer.

Lorenz, E. (1996). Predictability: A problem partially solved. In *Proceedings Seminar on Predictability*, pages 1–18. ECMWF: Reading Berkshire, UK.

Mandel, J., Cobb, L., and Beezley, J. D. (2011). On the convergence of the ensemble Kalman filter. *Applications of Mathematics*, 56(6):533–541.

Marin, J., Pudlo, P., Robert, C., and Ryder, R. (2012). Approximate Bayesian computational methods. *Statistics and Computing*, 22(6):1167–1180.

Mitchell, H. L. and Houtekamer, P. L. (2000). An adaptive ensemble Kalman filter. *Monthly Weather Review*, 128:416–433.

Myrseth, I. and Omre, H. (2010). Hierarchical ensemble Kalman filter. *SPE Journal*, 15(2):569–580.

Nott, D. J., Fan, Y., Marshall, L., and Sisson, S. A. (2014). Approximate Bayesian computation and Bayes linear analysis: toward high-dimensional ABC. *Journal of Computational and Graphical Statistics*, 23(1):65–86.

Nott, D. J., Marshall, L., and Ngoc, T. M. (2012). The ensemble Kalman filter is an ABC algorithm. *Statistics and Computing*, 22(6):1273–1276.

Ott, E., Hunt, B. R., Szunyogh, I., Zimin, A. V., Kostelich, E. J., Corazza, M., Kalnay, E., Patil, D. J., and Yorke, J. A. (2004). A local ensemble Kalman filter for atmospheric data assimilation. *Tellus*, 56A:415–428.

Pfanzagl, J. (1969). On the measurability and consistency of minimum contrast estimates. *Metrika*, 14:249–272.

Polson, N. G., Stroud, J. R., and Müller, P. (2008). Practical filtering with sequential parameter learning. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 70(2):413–428.

Pourahmadi, M. (2013). *High-Dimensional Covariance Estimation*. Wiley.

Raftery, A. E., Newton, M. A., Satagopan, J. M., and Krivitsky, P. N. (2007). Estimating the integrated likelihood via posterior simulation using the harmonic mean identity. In Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M., editors, *Bayesian Statistics 8*, pages 1–45. Oxford University Press.

Rao, V., Sandu, A., Ng, M., and Nino-Ruiz, E. (2015). Robust data assimilation using L_1 and Huber norms. *arXiv:1511.01593*.

Rauch, H. E., Striebel, C., and Tung, F. (1965). Maximum likelihood estimates of linear dynamic systems. *AIAA journal*, 3(8):1445–1450.

Rebeschini, P. and Van Handel, R. (2015). Can local particle filters beat the curse of dimensionality? *Annals of Applied Probability*, 25(5):2809–2866.

Robert, C. and Casella, G. (2004). *Monte Carlo Statistical Methods*. Springer.

Roh, S., Genton, M. G., Jun, M., Szunyogh, I., and Hoteit, I. (2013). Observation quality control with a robust ensemble Kalman filter. *Monthly Weather Review*, 141(12):4414–4428.

Rue, H., Martino, S., and Chopin, N. (2009). Approximate Bayesian inference for latent Gaussian models by using integrated nested Laplace approximations. *Journal of the Royal Statistical Society: Series B*, 71(2):319–392.

Schweppe, F. (1965). Evaluation of likelihood functions for gaussian signals. *IEEE transactions on Information Theory*, 11(1):61–70.

Shaby, B. A. (2014). The open-faced sandwich adjustment for MCMC using estimating functions. *Journal of Computational and Graphical Statistics*, 23(3):853–876.

Sigrist, F., Künsch, H. R., and Stahel, W. A. (2012). A dynamic nonstationary spatio-temporal model for short term prediction of precipitation. *Annals of Applied Statistics*, 6(4):1452–1477.

Sisson, S. A., Fan, Y., and Tanaka, M. M. (2007). Sequential Monte Carlo without likelihoods. *Proceedings of the National Academy of Sciences*, 104(6):1760–1765.

Snyder, C., Bengtsson, T., Bickel, P., and Anderson, J. L. (2008). Obstacles to high-dimensional particle filtering. *Monthly Weather Review*, 136:4629–4640.

Stroud, J. R., Katzfuss, M., and Wikle, C. K. (2016). A Bayesian adaptive ensemble Kalman filter for sequential state and parameter estimation. *arXiv:1611.03835*.

Stroud, J. R., Stein, M. L., Lesht, B. M., Schwab, D. J., and Beletsky, D. (2010). An ensemble Kalman filter and smoother for satellite data assimilation. *Journal of the American Statistical Association*, 105:978–990.

Szunyogh, I., Kostelich, E. J., Gyarmati, G., Kalnay, E., Hunt, B. R., Ott, E., Satterfield, E., and Yorke, J. A. (2008). A local ensemble transform Kalman filter data assimilation system for the NCEP global model. *Tellus*, 60A(1):113–130.

Tierney, L. and Kadane, J. B. (1986). Accurate approximations for posterior moments and marginal densities. *Journal of the American Statistical Association*, 81(393):82–86.

Tippett, M. K., Anderson, J. L., Bishop, C. H., Hamill, T. M., and Whitaker, J. S. (2003). Ensemble square-root filters. *Monthly Weather Review*, 131:1485–1490.

Tsyrulnikov, M. and Rakitko, A. (2015). Hierarchical Bayes ensemble Kalman filtering. *Physica D*.

Ueno, G. and Nakamura, N. (2014). Iterative algorithm for maximum-likelihood estimation of the observation-error covariance matrix for ensemble-based filters. *Quarterly Journal of the Royal Meteorological Society*, 140:295–315.

Ueno, G. and Nakamura, N. (2016). Bayesian estimation of observation error covariance matrix in ensemble-based filters. *Quarterly Journal of the Royal Meteorological Society*, 142:2055–2080.

Urbanek, S. and Rubner, Y. (2012). *emdist: Earth Mover's Distance*. R package version 0.3-1.

van Leeuwen, P. J. and Evensen, G. (1996). Data assimilation and inverse methods in terms of a probabilistic formulation. *Monthly Weather Review*, 124:2898–2913.

Wald, A. (1949). Note on the consistency of the maximum likelihood estimate. *The Annals of Mathematical Statistics*, 20(4):595–601.

Wang, X. and Bishop, C. H. (2003). A comparison of breeding and ensemble transform Kalman filter ensemble forecast schemes. *Journal of the Atmospheric Sciences*, 60:1140–1158.

West, M. (1987). On scale mixtures of normality. *Biometrika*, 74(3):694–697.

Whitaker, J. S., Hamill, T. M., Wei, X., Song, Y., and Toth, Z. (2008). Ensemble data assimilation with the NCEP global forecast system. *Monthly Weather Review*, 136(2003):463–482.

Wikle, C. K. (2002). A kernel-based spectral model for non-Gaussian spatial processes. *Statistical Modelling: An International Journal*.

Wood, S. N. (2010). Statistical inference for noisy nonlinear ecological dynamic systems. *Nature*, 466:1102–1104.